A synchronous rendering of hybrid systems for designing Plant-on-a-Chip (PoC)

Avinash Malik, Partha S Roop, Sidharta Andalam, Eugene Yip, and Mark Trew

Abstract—Hybrid systems are discrete controllers that are used for controlling a physical process (plant) exhibiting continuous dynamics. A hybrid automata (HA) is a well known and widely used formal model for the specification of such systems. While many methods exist for simulating hybrid automata, there are no known approaches for the automatic code generation from HA that are semantic preserving. If this were feasible, it would enable the design of a plant-on-a-chip (PoC) system that could be used for the *emulation* of the plant to validate discrete controllers. Such an approach would need to be mathematically sound and should not rely on numerical solvers. We propose a method of PoC design for plant emulation, not possible before. The approach restricts input/output (I/O) HA models using a set of criteria for well-formedness which are statically verified. Following verification, we use an abstraction based on a synchronous approach to facilitate code generation. This is feasible through a sound transformation to synchronous HA. We compare our method (the developed tool called Piha) to the widely used Simulink® simulation framework and show that our method is superior in both execution time and code size. Our approach to the PoC problem paves the way for the emulation of physical plants in diverse domains such as robotics, automation, medical devices, and intelligent transportation systems.

Index Terms—Hybrid Automata, Synchronous Languages, Plant-on-a-Chip (PoC), Code Generation.

arXiv:1510.04336v1 [cs.FL] 14 Oct 2015

1 Introduction

Cyber-physical Systems (CPS) [1], [2] encompass a wide range of systems where distributed embedded controllers are used for controlling physical processes. Examples range from the controller of a robotic arm, controllers in automotive such as drive-by-wire applications to pacemakers for a heart. Such systems, also called hybrid systems, combine a set of discrete controllers with associated physical plants that exhibit continuous dynamics. An individual plant combined with its controller is often formally described using Hybrid Automata (HA) [3], [2], [4]. These HA use a combination of discrete *locations* or *modes* and Ordinary Differential Equations (ODEs) to capture the continuous dynamics of an individual mode.

CPS pose many key challenges for their design, especially since they are inherently safety-critical. CPS must operate in a sound manner while preserving key requirements for both functional and timing correctness [5]. Moreover, there is a need for most CPS to be certified using functional safety standards such as IEC 61580 [6] (robotics and automation), ISO 26262 [7] (automotive), DO-178B [8] (flight control), US Food and Drug Administration (FDA) [9] (medical devices). Certification is a key barrier to rapid adoption of new technology as compliance costs are prohibitive. Also, in spite of products being certified, faults can occur and associated recalls are required, which can be immensely costly in terms of human lives and economics. For example, close to 200,000 pacemakers were recalled between 1990–2000 due to software related failures and this trend is continuing [10].

- *Authors are with the Department of Electrical and Computer Engineering, and Auckland Bioengineering Institute at The University of Auckland, New Zealand*
  E-mail: avinash.malik@auckland.ac.nz

Manuscript received xx xxxx xxxx

In spite of these safety crtical needs, current design practices often use tools without rigorous semantics. For example, simulation tools such as Matlab® Simulink® and Stateflow® [11], [12] are arguably the default standard in many fields, including automotive and medical. These tools enable the encoding of an HA using a set of Simulink® blocks to specify the continuous dynamics while the discrete aspects are modelled using Stateflow®. Due to the lack of rigorous semantics in Simulink®/Stateflow®, many problems have been reported in literature regarding the lack of good model fidelity. To overcome these problems, tools with rigorous semantics such as Ptolemy [13] and Zélus [12], [14] have been developed. These tools have developed mathematical semantics for the composition of the blocks and hence do not have semantic problems. However, all available tools, we are aware of, use numerical solvers during simulation to solve the ODEs. This poses two key challenges in the design of the CPS. Firstly, the simulation is inherently slow and potentially unsuitable for validating controllers in closed loop that can be critical to certification. A second problem is that simulation fidelity is dependent on the nature of the ODE solver and changing the solver may result in a different outcome.

Hardware-in-the loop simulation, also known as *emulation* [15], uses the actual plant in conjunction with the controller for validation. Emulation has been shown in many domains (such as motor controllers) to be an extremely effective method of validation. However, emulation of a hybrid system may not be feasible due to the following reasons. (1) The physical process may also be part of the design and not available for testing, e.g. a robotic arm and its controller are required but the actual robot may not exist yet. This is an even more significant hurdle for medical devices such as pacemakers, where emulating the actual system is often not feasible or may require animal

experiments that are challenging in terms of ethics, time and cost. (2) Simulation of physical process using existing tools may not be able to meet the timing requirements of emulation, e.g. real time responses.

Our work overcomes these barriers by proposing, for the first time, the concept of plant-on-a-chip (PoC). A PoC is an implementation of a hybrid system in C (and subsequently VHDL), so that a physical process can be emulated using a processor, ASIC, or FPGA. This device can then be used for closed-loop validation of controllers in wide ranging applications in CPS such as robotics, automotive, and medical devices.

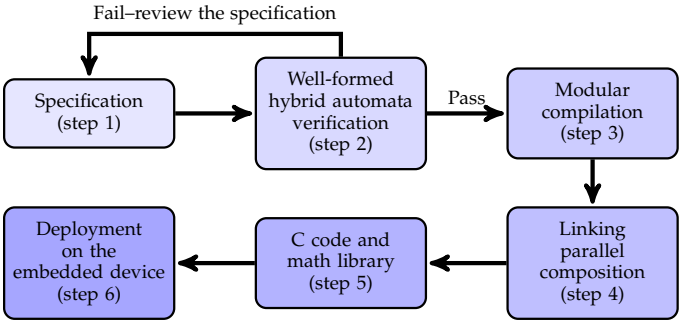
Fail–review the specification
Specification (step 1)
Well-formed hybrid automata verification (step 2)
Pass
Modular compilation (step 3)
Linking parallel composition (step 4)
C code and math library (step 5)
Deployment on the embedded device (step 6)

Fig. 1. An overview of the proposed methodology

The proposed approach significantly extends the synchronous approach [16] that is the default standard for designing safety critical systems in many domains including civil aviation. Consequently, our code generation relies on the well known tenets of the synchronous approach. An overview of the proposed methodology is presented in Figure 1. We start with a specification that describes a network of HA models of a CPS (step 1). To facilitate code generation, we constrain such models using a set of well formedness criteria, which can be statically verified (step 2). We then introduce an abstraction called synchronous hybrid automata (SHA), that provides the semantics of the generated code, and using this abstraction separate compilation can be performed for each HA (step 3). Such modular compilation is feasible due to the synchronous semantics of a network of well formed HAs (WHAs) presented in Section 3.4. Using this semantics, we perform a linking operation of the generated C code (step 4). Finally, the generated C code (step 5) can be compiled using a standard C compiler to deploy the binary on an embedded target platform (step 6). The proposed methodology is detailed in Section 4.

1.1 Related work

The design of CPS requires both formal analysis [17] and testing-based validation. Of particular importance is the need for emulation-based validation [15], as outlined above. Hybrid automata and their compositions [3], [4] provide a formal framework for the specification of hybrid systems. There are well known negative results on the decidability of the reachability problem over hybrid automata, i.e., even the simplest class, called *linear hybrid automata*, is undecidable [3]. Several restrictions have been proposed with associated semi-decidable results that are used for developing model checking-based solutions [18]. However, the negative decidability results mean that both formal analysis and modular code generation are critical [12].

Unlike the objective of modular code generation based on formal models, conventional CPS design relies on commercial tools such as Simulink® or Stateflow® [11]. However, Simulink® and Stateflow® lack formal semantics for the composition of the components leading to issues with simulation fidelity (see [12], [19] for details). In contrast to commercial tools, academic variants such as Ptolemy [13] and Zélus [12] are founded on formal models of computation. Ptolemy supports multiple models of computations and a specification can seamlessly combine them. This approach is extremely powerful for the simulation of complex CPS. These tools can model hybrid systems that use numerical solvers.

Sequential code generation from synchronous specifications is a very well developed area [16], [20]. For example, the code generator from SCADE specifications [21], [22] generates certified code according to the DO-178B standard. However, existing synchronous code generators are designed for discrete systems only and unsuitable for generating code from HA descriptions. Zélus [12], on the other hand, is an extension of the synchronous language Lustre [16] with ODEs for the design of hybrid systems. Zélus [12] overcomes the semantic ambiguities associated with Simulink®. While Ptolemy [13] and Zélus [12] are significant developments for CPS design due to their sound semantics, they are limited by their reliance on dynamic ODE solvers during simulation. Consequently, they are unsuitable for the emulation of physical processes. However, we note that approaches such as Ptolemy [13] and Zélus [12] are highly expressive as they can simulate any HA, while the proposed approach generates code from Well-formed Hybrid Automata (WHA) models, which are a subset of HA.

Code generation without using dynamic invocation of numerical solvers have been attempted in the past with limited success. Alur et al. generated code from statechart-like hierarchical HA models [23]. Like our approach, the authors proposed a requirement of restricting HA execution to what was statically verifiable. They enforced the requirement that "a given mode and the corresponding guard of any switch must overlap for a duration that is greater than the sampling period" [23]. However, we show that code generation is feasible without the need for such a restrictive approach. In a similar vein, Kim et al. also presented an approach for modular code generation from HA [24]. Their key idea was also based on the discretization of HA, requiring that "transitions are taken at time multiples of the discretization step." Also, a SystemC based code generation for HA has been proposed [25]. While the title claims code generation for HA, the actual algorithms are restricted to timed automata [26] and hence unsuitable for the general class of HA. This is not adequate for modelling the physical plant properties in many examples. In contrast, our approach (the developed tool called Piha[1]) generates code using a technique that

1. Piha is the name of a well known New Zealand surfing paradise. Piha (Pi-Hybrid Automata) also refers to synchronous HA and is not related to the concept of π in pi-calculus.

is similar in principle, but without these overly restrictive requirements. Hence, we demonstrate that many practical CPS, ranging from pacemakers to boiler systems, can be designed and implemented using the proposed approach.

1.2 Contributions and organisation

The key contributions of the paper are:

- *A software engineering (SE) approach for designing CPS*: We propose an approach that is based on key SE principles founded on formal methods and model-driven engineering [27]. The overall system design starts with the description of the system as a network of hybrid automata. We propose a new semantics of HA based on the notion of Well-formed Hybrid Automata (WHA). The WHA requirements can be statically checked and those that satisfy the WHA requirements are amenable to code generation.
- *Code generation without dynamic numerical solving*: We use a synchronous approach for the design of hybrid systems [12], [14]. Unlike these code generators, which rely on dynamic ODE solvers rendering them less suitable for PoC design, we generate code that uses no numerical solver at run-time. This is achieved by an algorithm that checks WHA requirements at compile time to create closed form solutions of ODEs using symbolic solvers. WHA based code generation is less restrictive compared with existing code generators for HA [23], [24] that also avoid dynamic ODE solvers.
- *Semantic preserving modular compilation of HA*: The proposed approach is based on a compositional semantics such that code can be generated for each HA separately.
- *A new PoC design approach*: We propose the first comprehensive approach for PoC design enabling the emulation of many CPS applications. Our approach, thus, paves the way for the reduction of the certification effort for safety-critical devices such as pacemakers [28], automotive electronics, and many other applications where certification is a requirement.
- *Practicality*: We have applied our approach to the design of several hybrid systems and compared our implementation with the widely used Simulink® and Stateflow®.
- *Software tool*: We have implemented these concepts in a software tool called Piha.

The paper is organised as follows. In section 2 we introduce HA as a modelling framework for CPS. We also define WHA as a restrictive class that is amenable for code generation. In section 3 we introduce an abstraction of WHA called synchronous hybrid automata and present its semantics. In section 4 we present algorithms for code generation based on the developed synchronous semantics. The algorithms include those to determine if an HA meets the requirements for code generation based on WHA. It also includes the back-end code generators from HA definitions to C using the developed synchronous semantics. Section 5 compares the proposed approach with existing tools such as Simulink® and Stateflow®. Finally in section 6 we make concluding remarks including the limitations of the developed approach and future research avenues.

2 Background: Hybrid automata

We use hybrid IO automata of Lynch et al. [29] to capture the behaviour of the hybrid system similar to the models adopted by Chen et al. [28]. This approach is specifically amenable to the modelling of both the plant and its adjoining controller, and hence ideal for the PoC design approach presented here. We refer to them as Hybrid Automata (HA) for convenience.

To illustrate the concept of HA, consider a water tank heating system, presented in Figure 2 (adapted from [3]), where the temperature of the water inside the tank is to be maintained between 20 and 100°C. The temperature is controlled by a digital controller that switches the heating system `ON` and `OFF` at specific times so as to maintain the temperature within the required range. The temperature of water inside a tank is modelled using Newton's law of heating and cooling, which can be represented as $x(t) = Ie^{-Kt} + h(1 - e^{-Kt})$. Here, I is the initial temperature, K is a constant that depends on the thermal conductivity of the tank, and h is a constant that depends on the power output of the heating system, such as a gas burner. An example trace of the temperature of water inside the tank for an initial temperature of 20°C is shown in Figure 2(b). The initial temperature is 20°C and the heating system is turned `ON`. Next, the temperature of water increases until it boils at 100°C, when the time instant is 13.2 seconds. Next, the temperature remains at 100°C until 17.8 seconds, when the heating system is turned `OFF` by the controller. The temperature of water drops until 24.6 seconds, when the controller switches `ON` the heating system again. We describe the water tank as a hybrid automata in Figure 2(c) and that of the gas burner in Figure 2(d).

The specification of the water tank heating system using an HA is shown in Figure 2(c). An HA captures the interaction between the plant (exhibiting *continuous dynamics*) and a controller (making *discrete mode switches*). The continuous evolution of some real-valued variables is modelled using a set of Ordinary Differential Equations (ODEs), which specify the rate of change of these variables. The controller effects the discrete mode switches and in each mode the plant exhibits different dynamics. As a consequence, the ODEs in the different modes (also called locations) may be different. For the water tank heating system, the HA can be in any one of four different locations t_1, t_2, t_3, and t_4, where t_1 is the initial location. The variable x is used to model the temperature of the water inside the tank. Each location has some flow predicates that specify the rate of change of the continuous variables. For example, in location t_2 the flow predicate that defines the rate of change of temperature when the heater is turned `ON` is specified as $\dot{x} = K(h - x)$. Invariants are associated with locations, e.g. $20 \leq x \leq 100$ is an invariant associated with location t_2. Execution remains in a location until the invariant holds or an egress transition triggers prior to this. Some locations may have initialization conditions that provide the initial values of the variables i.e. location t_1. A transition out of a locations is enabled when the input (event) is present and the jump condition associated with the transition holds. When a given transition is taken, the final value of the variables are updated. For example, the transition from t_2 to t_3 is enabled when the

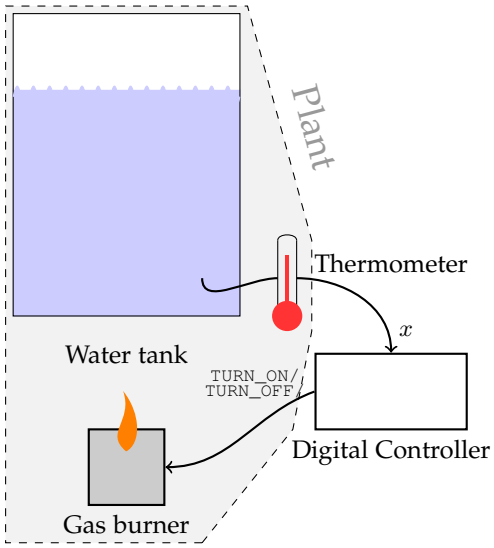

(a) Water tank system overview

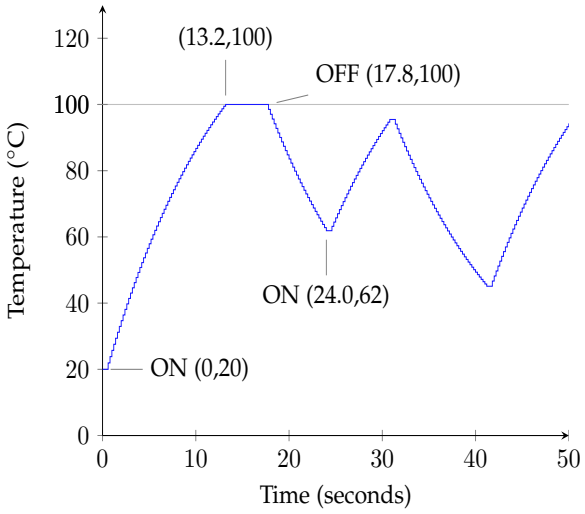

(b) A trace generated from our tool, called Piha

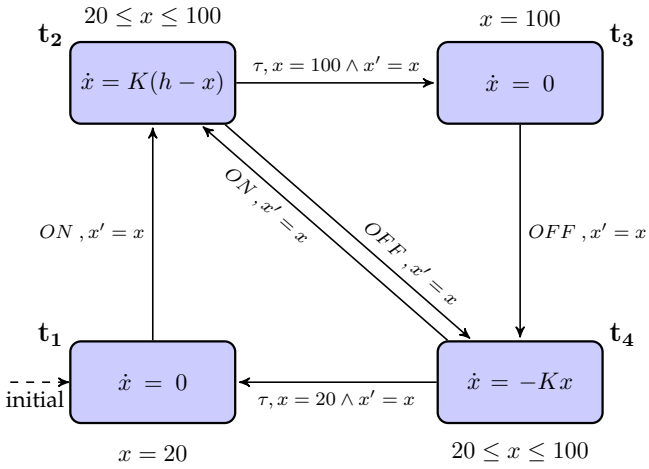

(c) Hybrid automata of the water tank

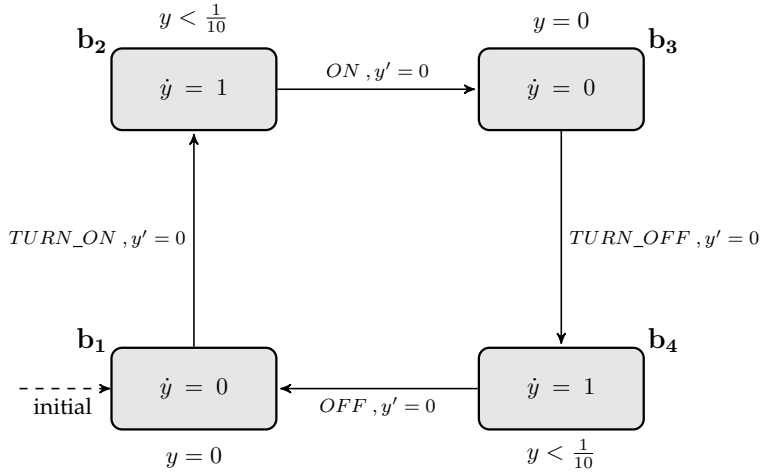

(d) Hybrid automata of the gas burner. The reaction time of the gas burner is $\frac{1}{10}$ seconds.

Fig. 2. An example of a water tank heating system – adapted from [3]. Figure 2(a) presents the plant and the discrete controller. Figures 2(c) and 2(d) captures the behaviour of the water tank and the gas burner. Figure 2(b) presents an example trace through a water tank heating system (generated from our tool). The trace is a set of discrete points sampled every 0.2 seconds.

value of temperature is $100°$C and as the transition is taken, the final value of x is updated using the equation $x' = x$, where x' is an update variable. The values of the update variables are used as the initial values of the variables in the destination location.

An HA [29] is defined using Definition 1. For the water tank example shown in Figure 2(c), $Loc = \{t_1, t_2, t_3, t_4\}$, $\Sigma = \{ON, OFF, \tau\}$, and an example edge in $Edge$ is (t_1, ON, t_2). There is a single continuous variable x representing the temperature of the water in the tank. Hence, $X = \{x\}$, $\dot{X} = \{\dot{x}\}$, and $X' = \{x'\}$. The only location that is marked initial is t_1 and $Init(t_1) : x = 20$. An example flow predicate is $Flow(t_1) : \dot{x} = 0$, which specifies that the temperature of the water inside the tank does not change. A jump predicate assigned to an edge specifies the condition needed to take a transition and the updating of values of some continuous variables when the transition is taken. For example, $Jump(t_4, \tau, t_1) : x = 20 \wedge x' = x$ specifies that the transition is taken when the value of the temperature is $20°$C and the updated value of the temperature is also $20°$C. We now formalise the HA using Definition 1.

Definition 1. *A hybrid automata is* $H = \langle Loc, Edge, \Sigma, Inv, Flow, Jump \rangle$ *where:*

- $Loc = \{l_1, .., l_n\}$ *representing* n *control modes or locations.*
- $Edge \subseteq Loc \times \Sigma \times Loc$ *are the set of edges between locations.*
- $\Sigma = \Sigma_I \cup \Sigma_O \cup \{\tau\}$ *is set of event names comprising of input, output, and internal events.*
- *Three sets for the set of continuous variables, their rate of change and their updated values represented as follows:* $X = \{x_1, .., x_m\}$, $\dot{X} = \{\dot{x_1}, .., \dot{x_m}\}$, *and* $X' = \{x'_1, .., x'_m\}$.
- $Init(l)$*: Is a predicate whose free variables are from* X*. It specifies the possible valuations of these when the HA starts in* l*.*
- $Inv(l)$*: Is a predicate whose free variables are from* X *and it constrains these when the HA resides in* l*.*

- *$Flow(l)$: Is a predicate whose free variables are from $X \cup \dot{X}$ and it specifies the rate of change of these variables when the HA resides in l.*
- *$Jump(e)$: Is a function that assigns to the edge 'e' a predicate whose free variables are from $X \cup X'$. This predicate specifies when the mode switch using 'e' is possible. It also specifies the updated values of the variables when this mode switch happens.*

The semantics of an HA is specified using the notion of timed transitions systems (TTS) using Definition 2.

Definition 2. *The semantics of an HA $H = \langle Loc, Edge, \Sigma, Inv, Flow, Jump \rangle$ is defined by a timed transition system $TTS = \langle Q, Q_0, \Sigma, \rightarrow \rangle$*

- *Q is of the form (l, v) where l is a location and $v \in [X \rightarrow \mathbb{R}]$ such that v satisfies $Inv(l)$. Q is called the state-space of H.*
- *$Q_0 \subseteq Q$ of the form (l, v) such that v satisfies $Init(l)$.*
- *$\rightarrow$ is the set of transitions consisting of either:*
  1) *Discrete transitions (instantaneous): For each edge $e = (l, \sigma, l')$, we have $(l, v) \xrightarrow{\sigma} (l', v')$ if $(l, v) \in Q$, $(l', v') \in Q$ and (v, v') satisfy $Jump(e)$.*
  2) *Continuous transition (delay): For each non-negative real δ, we have $(l, v) \xrightarrow{\delta} (l, v')$ if $(l, v) \in Q$, $(l, v') \in Q$, and there is a differentiable function $f : [0, \delta] \rightarrow \mathbb{R}^m$ (which is a witness function of the above transition where $m = |X|$) with first derivative $\dot{f} : (0, \delta) \rightarrow \mathbb{R}^m$ such that the following conditions hold:*
     - *$f(0) = v$ and $f(\delta) = v'$,*
     - *for all $\epsilon \in (0, \delta)$, $f(\epsilon)$ satisfies $Inv(l)$ and $(f(\epsilon), \dot{f}(\epsilon))$ satisfies $Flow(l)$.*

The TTS semantics of an HA encodes the state-space Q as all possible pairs of the form (l, v), where l is a location and v is a vector representing the valuation of the variables. In general, the state-space is infinite as there are infinite possible valuations of variables within any time-step δ. Transitions may be any one of two types. (1) Discrete transitions that lead to a mode switch and these transitions are *instantaneous*. This happens when the input event is present, the guard condition is true and the updates are performed as the transition is taken. (2) Continuous transitions that capture the passage of time when control resides in a location.

3 Semantic restrictions for efficient code generation

The HA model introduced in the previous section is very generic, expressive, and hence powerful. However, it has limitations for code generation. Since we want to generate code for emulating physical processes, i.e., for PoC design, the efficient, real-time response of the generated code is critical.

Another limitation of generic HA must also be carefully considered for code generation. As argued in the introduction, all known tools (including ones widely used in industry) solve ODEs at run-time using numeric techniques. Simulink® also introduces semantic ambiguities and there are problems with simulation fidelity, due to the dependence on the step size selected for the numerical solver [11]. For example, a numerical solver may fail to detect the peak of a continuous variable and may perform backtracking and other heuristic techniques, that affect the outcome of the simulation. As a consequence, robust tools, such as Ptolemy, are advocating the use of techniques such as quantized-state systems [30] as an alternative to dynamic numerical simulation.

Considering the difficulty posed by the above issues, we propose the well-formedness criteria for admitting a subclass of HA, that we term Well-formed Hybrid Automata (WHA), which is amenable to automatic code generation. Arbitrary predicates for describing invariants and guard conditions are replaced by more structured conditions inspired by clock constraints in timed automata [26]. We also move away from dynamic numerical solvers and instead use analytic solutions. Most importantly, we develop a synchronous semantics for the generated code that provides key advantages for modular compilation and sound transformations that are semantic preserving.

3.1 Well-formed hybrid automata (WHA)

A set is defined called $CV(X)$ that will be used for creating constraints over the continuous variables. This is provided in Definition 3. These constraints are either comparison of a variable with a natural number or Boolean combinations of these. Comparisons can only be performed using the following mathematical operators: $<, \leq, >, \geq$.

Definition 3. *Given a set of continuous variables X, then a set of constraints over these variables are defined as follows: $g := x < n | x \leq n | x > n | x \geq n | g \wedge g$ where, $n \in \mathbb{N}$ and $x \in X$. $CV(X)$ denotes the set of constraints over X. Without loss of generality we can extend these constraints to the domain of rational numbers $\mathbb{Q}$.*

A WHA must satisfy the following conditions and are defined in Definition 4.

- All invariants, jump and init predicates are members of the set $CV(X)$.
- In HA semantics, when control resides in a location, there are infinite valuations of variables in any given interval of time. This makes code generation difficult. To facilitate code generation, we make evaluations only at discrete intervals. These intervals correspond to the *ticks* of a synchronous program that will be used for code generation.
- The ODEs which define flow constraints in any location of the form $\dot{x} = f(x)$ must be *closed form* in nature. This property ensures that such ODEs are analytically solvable so that the witness functions needed for the generated code are analytically computable.
- All witness functions must be monotonic. This necessary property ensures that generated code can evaluate invariants and jump predicates using *saturation* rather than backtracking. This is detailed in Section 4.4.

We have defined WHA for both modular code generation and to enable formal verification, which is not the focus of the current paper. If only modular code generation is needed, some of the WHA restrictions may be relaxed further.

Definition 4. *A well-formed hybrid automata* $WHA = \langle Loc, Edge, \Sigma, Init, Inv, Flow, Jump \rangle$ *where*

- $Loc = \{l_1, .., l_n\}$ *representing* n *control modes or locations.*
- $\Sigma = \Sigma_I \cup \Sigma_O \cup \{\tau\}$ *is the input alphabet comprising of event names, including internal events.*
- $Edge \subseteq Loc \times \Sigma \times Loc$ *are the set of edges between locations.*
- *Three sets for the set of continuous variables, their rate of change and their updated values represented as follows:* $X = \{x_1, .., x_m\}$, $\dot{X} = \{\dot{x_1}, .., \dot{x_m}\}$, *and* $X' = \{x'_1, .., x'_m\}$.
- $Init(l)$: $x_1 = v_1, .., x_m = v_m$ *where* $v_1..v_m \in \mathbb{R}$.
- $Inv(l)$: $\in CV(X)$
- $Flow(l)$: *A set of ODEs that meet the closed form requirement. Also, the witness functions of the ODEs are monotonic.*
- $Jump(e)$: *This function maps each edge 'e' to the conjunction of a guard and an update. The guard is in* $CV(X)$ *and the update is in* $UP(X, X')$, *which specifies the value of the updated variables.*
- *Fairness: the system never remains in any location indefinitely, i.e., if a location invariant holds indefinitely, then a egress transition is enabled by the controller within a bounded time.*

The semantics of a WHA is also a TTS (Definition 2).

3.2 Background on the synchronous approach

The *synchronous approach* [16] is widely used for the management of concurrency and race conditions. It is, in particular, applied extensively to the design of safety-critical systems, such as the embedded software for the Airbus A320 which uses the SCADE language and associated tools [21].

The synchronous approach is ideal for the development of *reactive systems* [31]. A reactive system reacts to its environment continuously and, in order to remain reactive, the speed of the environment must be considered, i.e. outputs must be produced for the current set of inputs before the next set of inputs appear.

The synchronous paradigm assumes that the *idealised* reactive system produces outputs synchronously relative to inputs, i.e., the outputs are produced after zero delay. This is possible if the reactive system executes infinitely fast relative to its environment, which is known as the *synchrony hypothesis* [16]. Based on this hypothesis, time may be treated as a sequence of discrete instants or *ticks* with nothing happening between the completion of the current tick and the start of the next tick. This idea is prevalent in various fields of engineering such as digital logic design and control systems [32]. Typically, the simulation step size of the ODE solver used for the plant must match the sampling step size of the controller for correct system simulation [33] and this aspect is often difficult to achieve. Unlike this, in the synchronous approach, the notion of synchronous composition [27] formalises this automatically using compilation technology.

The synchrony hypothesis is valid so long as the minimum inter-arrival time of input events is longer than the maximum time required to perform a reaction or tick. This requires the computation of the worst-case reaction time (WCRT) of the synchronous program [34], [35], [36].

Synchronous languages, which are based on the synchrony hypothesis, include Esterel [37], Lustre [38] and Signal [39], where every program reaction occurs with respect to a logical clock that *tick*s. A new reaction starts at the beginning of a tick by taking a snapshot of the input signals, computing the outputs using the user specified logic, and emitting the output signals before the next tick starts. This is similar to the scan cycle of a programmable logic controller [27].

3.3 Synchronous HA (SHA)

Because hybrid systems extend reactive systems (which are usually discrete) with continuous dynamics, we need to make semantic adaptations to the synchronous model to facilitate code generation. We make the following assumptions:

- All ODEs can be solved using analytic methods to compute their witness functions.
- Execution is performed in discrete instants based on the synchronous approach and the duration of each instant is a fixed time δ.
- Due to the fairness assumption, the execution time spent in any location is always bounded.

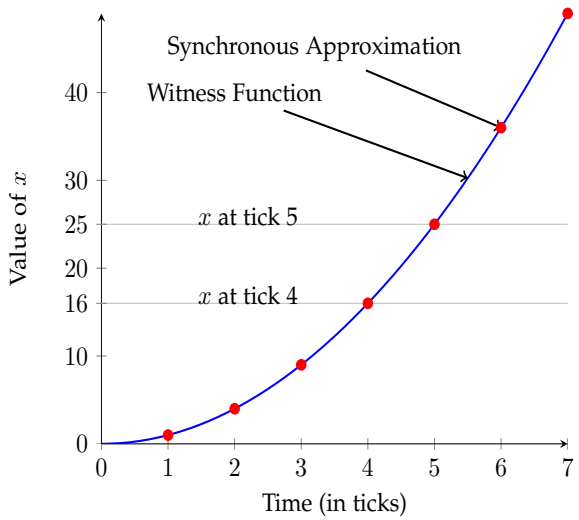

Fig. 3. An example visualisation of the synchronous approximation. Only observable at tick boundaries

Based on these assumptions, Definition 5 formalises the concept of SHA corresponding to any WHA. A SHA provides an *under-approximation* of the behaviour of a WHA (a subset of behaviours) due to the fact that in an SHA, the valuation of continuous variables are made at discrete instants and the value remains constant between two distinct valuations. This is visualised in Figure 3 and shows that the value of the continuous variable x is evaluated at discrete instants (or ticks) and remains unchanged between two instants (or ticks).

We also assume the following notation. Let $ValueInterval = \{[N_1, N_2] | N_1 \in \mathbb{N} \wedge N_2 \in \mathbb{N}\}$.

Definition 5. *Given a well-formed hybrid automata* $WHA = \langle Loc, Edge, \Sigma, Init, Inv, Flow, Jump \rangle$ *a SHA corresponding to this WHA is*

$SHA = \langle Loc, Edge, \Sigma, Init, Inv, Switness, Jump, Step, Nsteps, BoundryCond\rangle$ *where:*

- $Step = \delta \in \mathbb{R}^+$*: This specifies the duration of the synchronous instant.*
- $Nsteps : Loc \to \mathbb{N}$*: This specifies the maximum number of (logical) time steps that can be spent in any location.* $Nsteps(l) = k$ *implies that during any execution of the SHA the time spent in* l *must be less than or equal to* $k \times \delta$*.*
- $Switness : Loc \times \mathbb{N} \times \mathbb{R} \times \mathbb{R}^m \to \mathbb{R}^m$*: is the witness function that returns the valuation of all the continuous variables at any valid time step* k *in a given location* l*. It also takes as input the time step* δ *and the initial value from which execution begins in the location.*
- $BoundaryCond : (X \times Loc) \to ValueInterval$*: This function defines the interval in which the boundary value of any continuous variable lies in a given location. The boundary value of a continuous variable* x *is the value of the variable at time* 0*, when execution starts in that location, i.e.,* $x(0)$*.*

A SHA is an abstraction of the corresponding WHA. It inherits all the components of a WHA except that $Flow$ predicates are replaced by a composite witness function $Switness$. $Switness(l, k, \delta, i) = v_{k,l,i}$, which returns the evaluation of the witness function in location l at time step k. Here, $v_{k,l,i}$ is a vector representing the valuation of variables in X. We also denote $v_{k,l,i}(x) \in \mathbb{R}$ to represent the valuation of the continuous variable $x \in X$ in the k-th step in location l. We use the shorthand $x[k, l, i]$ to denote $v_{k,l,i}(x)$ and the shorthand $x[k]$ when the location l and the initial value of x in vector i, itself denoted as $i(x)$ is clear from the context.

We also introduce some other new components: $Step$, $Nsteps$, and $BoundaryCond$. $Step$ specifies the duration of a discrete instant or tick and $Nsteps$ maps every location to a natural number indicating the worst-case number of steps possible in that location during any execution of the SHA. $Step$ can be any value on the positive real-number line, usually obtained via *worst-case reaction time* (WCRT) analysis [34], [35], [36]. $Nsteps$ is computed statically using an algorithm presented in Section 4.2. $BoundaryCond(l, x)$ returns a closed interval of the form $[N_1, N_2]$, which means that the boundary value of $x(0)$ in location l is in $[N_1, N_2]$. This mapping is essential for computing the constants of integration and is explained in Section 4.2.

The semantics of a Synchronous Hybrid Automata (SHA) is provided as a Discrete Time Transition System (DTTS) in Definition 6. We assume that all transitions of a DTTS trigger relative to the ticks of the logical clock of the synchronous program.

Definition 6. *The semantics of a* $SHA = \langle Loc, Edge, \Sigma, Init, Inv, Switness, Jump, Step, Nsteps, BoundaryCond\rangle$ *is a* $DTTS = \langle Q, Q^0, \Sigma, \to\rangle$ *where*

- *The state-space is* Q*, where any state is of the form* (l, v, i, k) *where* l *is a location,* i *is the initial valuation of the variables when execution begins in the location and* v *is the valuation at the* k*-th instant.*
- $Q^0 \subseteq Q$ *where every* $q^0 \in Q^0$ *is of the form* (l, v^0, i, k) *such that* v^0 *satisfies* $Init(l)$*.*
- *Transitions are of two types:*
  - Inter-location transitions *that lead to mode switches: These are of the form* $(l, v, i, k) \xrightarrow{\sigma} (l', v', i', 0)$ *if* $(l, v, i, k) \in Q$*,* $(l', v', i', 0) \in Q$*,* $e = (l \xrightarrow{\sigma} l') \in Edge$ *and* (v, v') *satisfy* $Jump(e)$*.*
  - Intra-location transitions *made during the execution in a given mode / location: These are of the form* $(l, v, i, k) \to (l, v', i, k + 1)$ *if* $(l, v, i, k) \in Q$*,* $(l, v', i, k + 1) \in Q$*,* (v, v') *satisfy* $Inv(l)$*,* $Switness(l, k, \delta, i) = v$ *and* $Switness(l, k + 1, \delta, i) = v'$*.*

3.4 Composition

The developed synchronous semantics enables modular code generation. We can compile each WHA separately (step 3) in Figure 1 and then the generated codes are linked together in step 4, which is similar to the linking process in conventional compilation. We formalise the synchronous parallel composition of SHAs using Definition 7. This facilitates the seamless linking process, described in Section 4.5.

Definition 7. *Given* $SHA_1 = \langle Loc_1, Edge_1, \Sigma_1, Init_1, Inv_1, Switness_1, Jump_1, Step_1, Nsteps_1, BoundaryCond_1\rangle$ *and its semantics* $DTTS_1 = \langle Q_1, Q_1^0, \Sigma, \to_1\rangle$ *and* $SHA_2 = \langle Loc_2, Edge_2, \Sigma_2, Init_2, Inv_2, Switness_2, Jump_2, Step_2, Nsteps_2, BoundaryCond_2\rangle$ *and its semantics* $DTTS_2 = \langle Q_2, Q_2^0, \Sigma_2, \to_2\rangle$ *and* $Shared_v = X_1 \bigcap X_2$ *and* $Shared_e = \Sigma_1 \bigcap \Sigma_2$ *denoting a set of shared variables and events, respectively, which satisfy the following conditions:*

1) *All writes to shared variables* $x \in Shared_v$ *or the emission of shared events* $e \in Shared_e$ *are mutually exclusive.*
2) *Any read to a shared variable* $x \in Shared_v$ *or a shared event* $e \in Shared_e$ *accesses the value of the variable / event in the previous tick, denoted* $pre(x)/pre(e)$*.*

$DTTS_1 || DTTS_2 = DTTS\langle Q, Q^0, \Sigma, \to\rangle$ *where:*

- *The state-space is* $Q \subseteq Q_1 \times Q_2$*.*
- $Q^0 \subseteq Q_1^0 \times Q_2^0$*.*
- *Transitions* $\to$ *are of two types:*
  - Inter-location transitions *of the form:*

    (`Rule Inter-Inter`*)* $(q_1, q_2) \xrightarrow{\sigma_1 \wedge \sigma_2} (q_1', q_2')$ *where* $q_1 = (l_1, v_1, i_1, k)$*,* $q_2 = (l_2, v_2, i_2, k)$*,* $q_1' = (l_1', v_1', i_1', 0)$*,* $q_2' = (l_2', v_2', i_2', 0)$*.* $e_1 = (q_1 \xrightarrow{\sigma_1} q_1') \in Edge_1$ *and* (v_1, v_1') *satisfy* $Jump(e_1)$*.* $e_2 = (q_2 \xrightarrow{\sigma_2} q_2') \in Edge_2$ *and* (v_2, v_2') *satisfy* $Jump(e_2)$*.*

    (`Rule Inter-Intra`*)* $(q_1, q_2) \xrightarrow{\sigma_1} (q_1', q_2')$ *where* $q_1 = (l_1, v_1, i_1, k)$*,* $q_2 = (l_2, v_2, i_2, k)$*,* $q_1' = (l_1', v_1', i_1', 0)$*,* $q_2' = (l_2, v_2', i_2', 0)$*.* $e_1 = (q_1 \xrightarrow{\sigma_1} q_1') \in Edge_1$ *and* (v_1, v_1') *satisfy* $Jump(e_1)$*.* $q_2, q_2' \in Q_2$ *and* (v_2, v_2') *satisfy* $Inv(l_2)$*. Finally,* $i_2' = v_2$ *and* $v_2' = v_2$*.*

    (`Rule Intra-Inter`*)* $(q_1, q_2) \xrightarrow{\sigma_2} (q_1', q_2')$ *where* $q_1 = (l_1, v_1, i_1, k)$*,* $q_2 = (l_2, v_2, i_2, k)$*,* $q_1' = (l_1, v_1', i_1', 0)$*,* $q_2' = (l_2', v_2', i_2', 0)$*.* $e_2 = (q_2 \xrightarrow{\sigma_2} q_2') \in Edge_2$ *and* (v_2, v_2') *satisfy* $Jump(e_2)$*.* $q_1, q_1' \in Q_1$ *and* (v_1, v_1') *satisfy* $Inv(l_1)$*. Finally,* $i_1' = v_1$ *and* $v_1' = v_1$*.*

  - Intra-location transitions *of the form:*

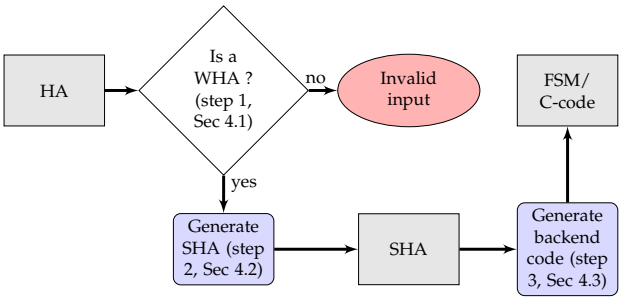

Fig. 4. Process for compiling a single Hybrid Automata (HA)

(`Rule Intra-Intra`) $(q_1, q_2) \rightarrow (q_1', q_2')$ *where* $q_1 = (l_1, v_1, i_1, k)$, $q_2 = (l_2, v_2, i_2, k)$, $q_1' = (l_1, v_1', i_1, k+1)$, $q_2' = (l_2, v_2', i_2, k+1)$, *and* $q_1, q_1' \in Q_1$, (v_1, v_1') *satisfy* $Inv(l_1)$, $Switness(l_1, k, \delta, i_1) = v_1$ *and* $Switness(l_1, k+1, \delta, i_1) = v_1'$. *Similarly,* $q_2, q_2' \in Q_2$, (v_2, v_2') *satisfy* $Inv(l_2)$, $Switness(l_2, k, \delta, i_2) = v_2$ *and* $Switness(l_2, k+1, \delta, i_2) = v_2'$ *for any* δ.

In Definition 7, the DTTS corresponding to two SHAs composed in parallel is computed by composing their respective DTTSs. The state-space Q of the resultant DTTS is a subset of the product of the state-space of the individual constituents. The initial state-space Q^0 is also a subset of the product of the initial state-space of the constituents. The transition relation consists of two types of transitions. The inter-location transitions happen when any one of the constituents or both constituents make an inter-location transition (there are three different possibilities). `Rule Inter-Inter` states that both constituents can take a discrete transition to new locations.

`Rule Inter-Intra` states that when the first constituent Q_1 takes an inter-location transition from location l_1 to l_1', Q_2 is also *forced* to make such a transition. But, the resultant location of Q_2 does not change, i.e., Q_2 can only take a transition from l_2 to l_2. Moreover, upon taking the transition the initial value i_2 is set to the current valuation v_2, consequently implying that $v_2' = v_2$. Once in the new state (l_1', l_2) vectors v_1 and v_2 start evolving according to their individual witness functions. `Rule Intra-Inter` is the dual of `Rule Inter-Intra`. Finally, intra-location transition in the composition (`Rule Intra-Intra`) happens only when both constituents make an intra-location transition.

4 Methodology for Code Generation

Figure 4 outlines the approach used to compile a single HA. The compilation process consists of three steps. We describe all three steps, in Sections 4.1-4.3, using the water tank HA, which has been reproduced in Figure 5(a).

4.1 Static analysis of hybrid automata

Given an HA, the first step (Figure 4) in the compilation process is to statically determine if the well-formedness criteria defined in Section 3.1 is respected. If the HA does not respect the well-formedness criteria, then an error is generated.

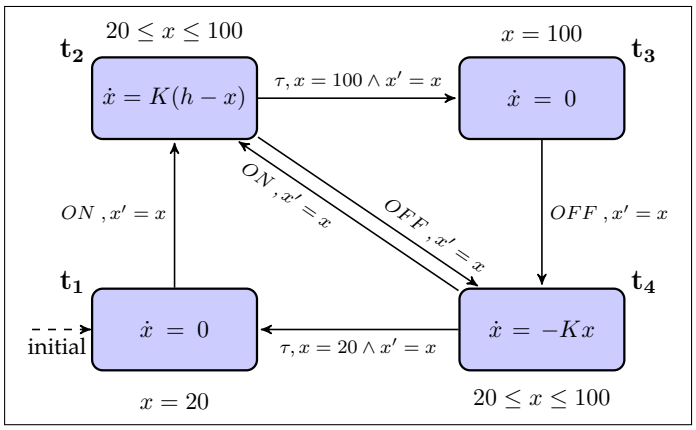

(a) Hybrid Automata (HA), see Definition 1. Reproduced from Figure 2(c). Symbols K and h are constants with values 0.075 and 150, respectively.

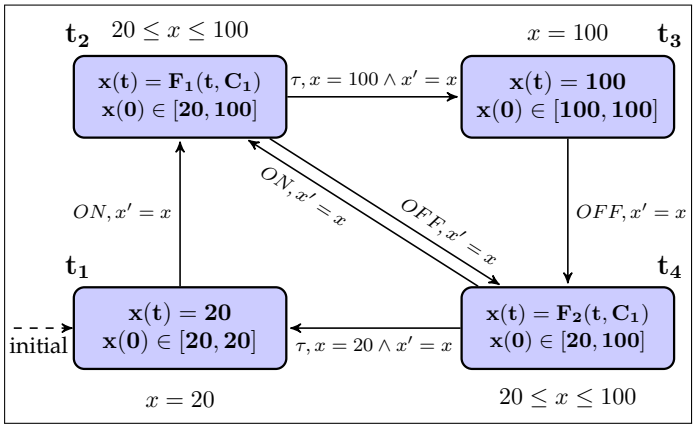

(b) Synchronous Hybrid Automata (SHA), see Definition 5. Flow predicates are described using witness functions. Values of the constants h and K are 150 and 0.075, respectively. Furthermore, witness function $F_1 = C_1 e^{-0.075 \times t} + 150.0$ and $F_2 = C_1 e^{-0.075 \times t}$

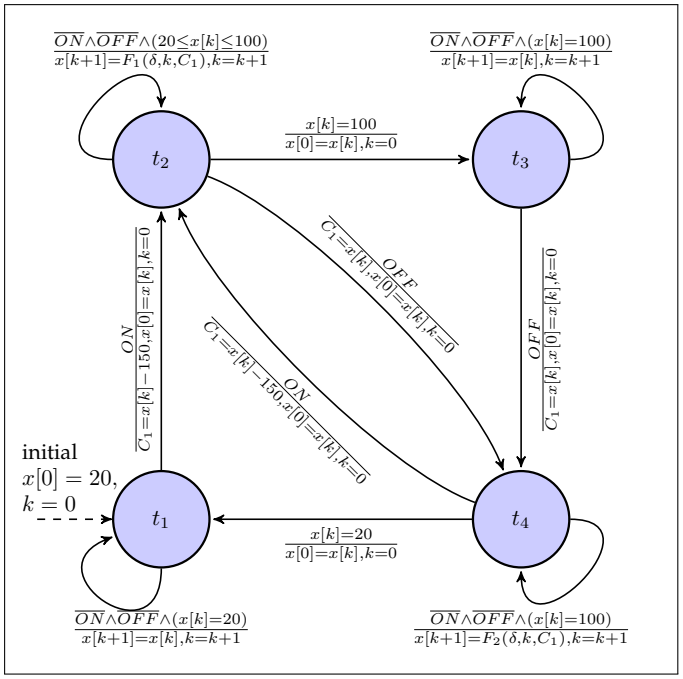
$\frac{\overline{ON} \wedge \overline{OFF} \wedge (20 \leq x[k] \leq 100)}{x[k+1] = F_1(\delta, k, C_1), k = k+1}$
$\frac{\overline{ON} \wedge \overline{OFF} \wedge (x[k] = 100)}{x[k+1] = x[k], k = k+1}$
t_2
$\frac{x[k] = 100}{x[0] = x[k], k = 0}$
t_3
$\frac{ON}{C_1 = x[k] - 150, x[0] = x[k], k = 0}$
$\frac{OFF}{C_1 = x[k], x[0] = x[k], k = 0}$
$\frac{ON}{C_1 = x[k] - 150, x[0] = x[k], k = 0}$
$\frac{OFF}{C_1 = x[k], x[0] = x[k], k = 0}$
initial
$x[0] = 20,$
$k = 0$
t_1
t_4
$\frac{x[k] = 20}{x[0] = x[k], k = 0}$
$\frac{\overline{ON} \wedge \overline{OFF} \wedge (x[k] = 20)}{x[k+1] = x[k], k = k+1}$
$\frac{\overline{ON} \wedge \overline{OFF} \wedge (x[k] = 100)}{x[k+1] = F_2(\delta, k, C_1), k = k+1}$

(c) Synchronous Witness Automata (SWA). We abuse the notation $x[k]$ to update the value of x, although $x[k]$ represents the valuation of the continuous variable x. This physical time $t = k \times \delta$, where k is the logical tick and δ is the tick length.

Fig. 5. The water tank component from the running example

Algorithm 1 The algorithm to check the well-formedness criteria of a HA

Input: HA ha
Output: Boolean

```
1: gset ← ∅
2: for all edges ∈ ha do
3:    for all guards ∈ edges do
4:       gset ← gset ∪ guards
5:    end for
6: end for
7: for all loc ∈ ha do
8:    for all invs ∈ loc do
9:       gset ← gset ∪ invs
10:   end for
11: end for
12: // check that all location invariants and jump conditions are of type CV(X)
13: for all g ∈ gset do
14:   assert type(g) ∈ CV(X)
15: end for
16: for all loc ∈ ha do
17:   for all ode ∈ loc do
18:      if solve_ode(ode) then
19:         (R1, R2) ← solve(ode.rhs > 0) [2]
20:         (R1', R2') ← solve(ode.rhs < 0)
21:         [N1, N2] ← get_inv_bounds(ode, loc)
22:         if (R1, R2) ∩ [N1, N2] ≠ ∅ ∧ (R1', R2') ∩ [N1, N2] ≠ ∅ then
23:            // Slope of witness function changes sign
24:            throw Exception ("Not a WHA")
25:         else
26:            return True
27:         end if
28:      else
29:         // No closed form solution, hence not a WHA
30:         throw Exception("Not a WHA")
31:      end if
32:   end for
33: end for
```

The procedure to verify the well-formedness criteria is presented in Algorithm 1, which takes an HA as input. Three well-formedness criteria need to be guaranteed. First, the invariants and the jump conditions need to be of the form $CV(X)$ (Definition 3). Lines 2-15 guarantee that this criterion is met. The second and third criteria require that each ODE, in every location, of the HA should have a closed form solution and should be monotonic. Lines 16-33 ensure that these criteria are satisfied.

Consider the running example HA – the water tank system presented in Figure 5(a). Lines 2-15 in Algorithm 1 collect all the invariant and jump conditions from the locations and the edges, respectively. Once collected in set $gset$, an assertion statement guarantees that all these conditions are of the form $CV(X)$ (line 14). Lines 16-33 iterate through each location of the HA. Upon visiting a location, all ODEs within the location are solved symbolically (line 18). If no closed form solution exists, then an exception is generated (line 30).

Given that a closed form solution exists, we then guarantee that all ODEs in a location are monotonic (not necessarily strictly monotonic). We use the definition that any given (witness) function is considered monotonic if and only if the first derivative of the function does not change sign [40]. Flow conditions evolve one or more ODEs within a given location as long as the invariant on the location is not violated. Hence, in our case, the definition of a monotonic function can be made more specific: any given (witness) function is monotonic if and only if its first derivative does not change sign within the interval specified by the invariant(s) of the location.

The right hand side of the ODEs specify the first derivatives of the witness functions. We need to ensure that the right hand side expression of the ODE (the slope) does not change signs within the invariant bounds. The lines 19-24 ensure that these conditions are satisfied. Line 19 obtains the real number line interval (denoted by (R_1, R_2)) such that the derivative of the witness function is always greater than zero, i.e., an increasing function. Line 20 obtains the real number line interval (denoted by (R_1', R_2')), such that the first derivative of the witness function is less than zero. Next, we obtain the invariant interval (denoted $[N_1, N_2]$), bounding the value of the evolving variable in the ODE, from the invariant(s) on the location. A non-empty interval $(R_1, R_2) \cap [N_1, N_2]$ indicates that the witness function is increasing within the location intervals. Similarly, a non-empty interval $(R_1', R_2') \cap [N_1, N_2]$ indicates that the witness function is a decreasing function within the location invariants. If both sets are non-empty, the witness function increases and decreases within the invariants specified on the location, and hence, the witness function is not monotonic.

Consider the running example in Figure 5(a) and specifically consider location t_2. The steps to determine if the witness function in location t_2 is monotonic are as follows:

1) *Compute the real number line interval for x such that the right hand side of the ODE is strictly greater than zero*: $x \in solve((0.075 * (150 - x) > 0)$ results in $x \in (-\infty, 150)$
2) *Compute the real number line interval for x such that the right hand side of the ODE is strictly less than zero*: $x \in solve((0.075 * (150 - x) < 0)$ results in $x \in (150, \infty)$
3) *Obtain the invariant bounds on x*: For location t_2, this interval is obtained from invariant: $20 \leq x \leq 100$. Hence, the invariant bounds are: $x \in [20, 100]$.
4) *Check if the witness function is monotonic*: The intersection $(-\infty, 150) \cap [20, 100] = [20, 100]$, but $(150, \infty) \cap [20, 100] = \emptyset$, and hence, the witness function in location t_2 is an increasing (and monotonic) function.

For the running example in Figure 5(a), all the well-formedness criteria are satisfied and hence, the water tank system is a WHA.

4.2 Generation of SHA

This section describes step 2 in the compilation procedure from Figure 4. Once a given HA is is determined to be a WHA, an SHA is generated from the WHA. This procedure translates all ODEs in every location of the WHA into their

2. $solve$ is a polymorphic function from Sympy (a python library) that can be used to solve many different types of problems. We make use of the polymorphic nature of $solve$ in Algorithms 1 and 2

Algorithm 2 The algorithm to generate a SHA from a WHA

Input: HA ha
Output: SHA sha

```
1: for all loc ∈ ha do
2:   for all ode ∈ loc do
3:     times ← ∅
4:     eq ← solve_ode(ode)
5:     ode ← eq
6:     if sign(diff(eq, t)) > 0 then
7:       times ← times ∪ solve(eq, min(Init), max(Inv))
8:     else if sign(diff(eq, t)) < 0 then
9:       times ← times ∪ solve(eq, max(Init), min(Inv))
10:    else
11:      raise Warning ("Fairness required")
12:      times ← ∅
13:      break
14:    end if
15:  end for
16:  if min(times) = ∞ then
17:    raise Warning ("Fairness required")
18:  end if
19: end for
20: return this
```

closed form solutions (witness functions). It also computes the worst-case bound $Nsteps$ (if one exists) as defined in Definition 5.

The procedure to generate the SHA from a WHA is presented in Algorithm 2. The algorithm visits every location in the WHA. For each ODE in the location, a witness function is obtained (line 4). Next, the algorithm attempts to determine the maximum time that the HA will spend in every location (lines 6-18) in the worst-case. The algorithm first detects the slope of each ODE (*ode* in Algorithm 2) for any given location (*loc* in Algorithm 2), by differentiating the witness function with respect to time t (line 6). If the slope is increasing (line 6) then the witness function (eq in Algorithm 2) is solved for t. We do the reverse for an ODE with decreasing slope (line 9). If the slope is constant, in case of a non-changing ODE, a warning is raised stating that $egress$ (fairness) condition is needed to make progress out of the location (line 11). Given the various worst-case times in the set $times$, for all ODEs in a location, the minimum value amongst $times$ needs to be bounded to guarantee at compile time that $Nsteps$ is bounded, otherwise an egress condition is needed to make progress out of the location (line 17).

Consider the example of the SHA shown in Figure 5(b). The algorithm visits each location in the WHA (Figure 5(a)) and replaces each ODE with its equivalent closed form solution (Algorithm 2, line 5). The closed form solution of all ODEs for the water tank is shown in Figure 5(b), given $h = 150$ and $K = 0.075$, respectively.

Consider location t_1 in Figure 5(b). Line 6, in Algorithm 2, differentiates the witness function with respect to time (t), resulting in a sign of 0, indicating a non-changing witness function. Hence, a warning that an $egress$ transition is needed to make progress out of location t_1 is raised. This is expected, because as seen in Figure 5(a), the ODE $\dot{x} = 0$ does not change the value of x and the invariant $x = 20$ holds, hence progress out of location t_1 is only possible upon reception of ($egress$) event `ON`.

Similarly, differentiating the witness function with respect to time in location t_4, gives the solution $-0.075 \times C1$, where $C1$ is the constant of integration. In our framework, the programmer annotates every location with the possible range of (initial) values that the continuous variable(s) might take upon entering the location. For example, $BoundaryCond(t_4, x)$ is denoted as $x(0) \in [20, 100]$ for location t_4[3]. Given that the initial value interval is positive, the slope of the witness function is detected to be negative (since C1 can only take a positive value) and hence, the witness function in location t_4 is a decreasing function. Once this is deduced, the worst-case time spent in location t_4 is computed at line 9, in Algorithm 2, by substituting the initial value of x as the maximum value from the *ValueInternal*, i.e., $x(0) = max(20, 100) = 100$ and the final value of x as the minimum value from amongst the invariant on x in location t_4, i.e., $x(t) = min(20, 100) = 20$.

Given that the witness function $x(t)$ is a decreasing function with an initial value of 100, we compute, using $solve$ on line 9, the time t it takes for function $x(t)$ to reach its final value 20. This time t, if it can be computed, gives the maximum time that the system will remain in location t_4 before making progress to a new location. The *solve* procedure proceeds as follows:

1) *Compute the value of the constant of integration*: For location t_4, $x(t) = C_1 \times e^{-0.075 \times t}$ as shown in Figure 5(b). First of all, we compute the constant of integration, using the initial value $x(0) = 100 \Rightarrow 100 = C_1 \times e^{-0.075 \times 0}$. Therefore, $C_1 = 100$ and hence $x(t) = 100 \times e^{-0.075 \times t}$.
2) *Computing time t*: Once, we have computed the constant of integration, we can easily compute the time t required for $x(t)$ to decrease from 100 to 20 as follows: $x(t) = 20 \Rightarrow 20$. Therefore, $20 = 100 \times e^{-0.075 \times t}$. Hence $t = \frac{ln(\frac{20}{100})}{-0.075}, t \approx 214.6$.

The system will make progress out of location t_4, in the worst-case, after approximately, 214.6 units of time. Similarly, we can bound the worst-case time spent in a location with increasing, rather than decreasing witness functions. In case a location consists of multiple witness functions, the worst-case time spent in that location is the minimum amongst the worst-case times computed for all the witness functions in that location. If the minimum amongst all the worst-case times is ∞, then a warning stating that an egress transition is necessary to make progress out of a location is generated as seen on line 17 in Algorithm 2.

4.3 Backend code generation

Finally, we describe the last step, step 3, in the compilation procedure presented in Figure 4. After the verification of WHA requirements and the generation of the SHA, backend code is generated. In the backend, the SHA is represented as a *Synchronous Witness Automata (SWA)*, which captures the behaviour of the SHA as a synchronous state machine. The SWA is a discrete variant which, when executed, produces

3. The possible initial value intervals can be computed automatically from the HA, but describing this is out of the scope of this paper.

the desired behaviour as a DTTS (Definition 6). During the execution, a task known as *saturation* may also be needed while taking discrete transitions. This is formalised in Section 4.4. The SWA is formalised using Definition 9. We start by defining $CV(X[k])$ in Definition 8 that is essential for the definition of the transition relation of an SWA.

Definition 8. *Let $x[k]$ denote the k-th ($k \in N$) valuation of any variable $x \in X$ and let $X[k]$ denote the set of all such k-valuations of variables in X.*

We denote $CV(X[k])$ the set of constraints over $X[k]$: $g := x[k] < n | x[k] \leq n | x[k] > n | x[k] \geq n | g \wedge g$ where, $n \in \mathbb{N}$.

Definition 9. *A synchronous witness automata (SWA) $\mathcal{S} = \langle S, S_0, \Sigma, X, Updates, \rightarrow \rangle$, which corresponds to the $SHA = \langle Loc, Edge, \Sigma, Init, Inv, Switness, Jump, Step, Nsteps, BoundaryCond \rangle$ where*

- *S denotes the set of states of $\mathcal{S}$ and $S = Loc$.*
- *S_o is a subset of initial states.*
- *Σ is a set of events.*
- *X is a set of continuous variables*
- *$Updates$ represent a set of updates. A given update may capture the emission of an output event, the update of a state variable in a given tick using its witness function or the initialization of an integration constant or the time instant.*
- *$\rightarrow \subseteq S \times \mathcal{B}(\Sigma) \times CV(X[k]) \times 2^{Updates} \times S$ represents the transition relation. Here $\mathcal{B}(\Sigma)$ denotes the set of Boolean formulas over Σ.*

We use the shorthand, $(s, \frac{e \wedge g}{u}, s')$ to represent transitions. Here, $e \in \mathcal{B}(\Sigma)$, $g \in CV(X[k])$ and $u \in 2^{Updates}$.

The generated C-code is an SWA. The SWA for the water tank generated from the SHA (Figure 5(b)) is shown in Figure 5(c). Every location in the SHA has an equivalent state in the SWA, as shown in Figure 5(c). Furthermore, the witness functions in the locations of the SHA are moved to transitions in the SWA. For example, the SWA self-transition from state t_2 to t_2 represents the evolution of the continuous variable x in discrete steps $k \in [0, Nsteps)$. In the SWA of Figure 5(c), we label the transitions with the form: $\frac{antecedent}{consequent}$. The self-transition from state t_2 to state t_2 is labeled as: $\frac{\overline{ON} \wedge \overline{OFF} \wedge (20 \leq x[k] \leq 100)}{x[k+1] = F_1(\delta, k, C_1), k = k+1}$, which states that while the value of x at tick k is between 20 and 100 *and* no events (`ON` or `OFF`) are detected, x evolves to the new value depending upon the witness function F_1. C_1 is the constant of integration. The transition guard has a one-to-one correspondence with the invariant condition on location t_2 in Figure 5(b). We add self-transitions to all states of the generated SWA, which evolve all the continuous variables in the corresponding location in the SHA. The generated SWA also consists of all the edges from the SHA, e.g., transition $(t_1, \frac{ON \wedge x[k]=20}{C_1 = x[k] - 150, x[0] = x[k], k=0}, t_2)$, which corresponds to the jump and updates on the transition between locations t_1 and t_2 in the corresponding SHA in Figure 5(b).

The C-code representing the SWA in Figure 5(c) is shown in Figure 6. We have only shown the states t_1 and t_4 along with their associated transitions and functions. There is a literal one-to-one correspondence between the SWA in Figure 5(c) and the C-code in Figure 6. We use the variables x and x_u to denote the $x[k]$ and $x[k+1]$ valuations of the

```
double t4_ode_1(double d, double k, double C1){
   return C1*exp(-0.075*d*k);
}
double t4_init_1(double x_u){
   return = x_u;
}
enum states watertankR(enum states cstate,
                       enum states pstate){
  switch (cstate){
  case (t1):
    if(x == 20 && !ON && !OFF){
      if (pstate != cstate) C1 = t1_init_1(x);
      x_u = t1_ode_1(C1);
      k=k+1;
      cstate = t1;
    }
    else if(!OFF){
      k=1; cstate=t2; x_u = x;
    }
    break;
  case (t2): ...
  case (t3): ...
  case (t4):
    if(x >= 20 && x <= 100 && !ON && !OFF){
      if (pstate != cstate) C1 = t4_init_1(x);
      x_u = t4_ode_1(d, k, C1);
      k=k+1;
      cstate = t4;
    }
    else if(ON){
      k=0; cstate=t2; x_u = x;
    }
    else if(x == 20){
      k=0; cstate=t1; x_u = x;
    }
    break;
  }
  return cstate;
}
```

Fig. 6. C-code generated for the SHA in Figure 5(c)

continuous variable x, respectively. The witness functions are represented as functions in C (e.g., lines 1-10 in Figure 6).

The witness functions are incomplete, in the sense that the value of the constant of integration (e.g., $C1$ on line 3) needs to be computed. Computing this constant of integration is equivalent to solving the *initial value problem*. It is well known that the value of the constant of integration, in the witness function, depends upon the initial value of the witness function. The initial value, at time 0, of any witness function is either: (1) the specified initial value of the continuous variable that the witness function updates, as shown in Figure 5(c) with the dashed arrow, or (2) the updated value of the continuous variable on an `Inter` location transition.

Consider the two transitions, one from location t_2 to t_4 and the other from t_4 to t_2 in Figure 5(c). Variable x is updated by the witness function F_2 on the self-transition in state t_4. Function F_2 takes as input arguments the current

tick number, $k \in \mathbb{N}$, the step size of the tick, $\delta \in \mathbb{R}^+$, and the constant of integration. The argument k starts from 0 and δ is a constant. We determine the value of the constant of integration when entering state t_4, by making C_1 the subject. Since $F_2 = C_1 \times e^{\delta \times k \times -0.075}$, we set $k = 0$ and making C_1 the subject gives $C_1 = F_2(0)$, i.e., C_1 takes the value of $x[0]$, which as we know from Figure 5(b) is the final value of x in the previous state, since $x' = x$. Hence, we get $C_1 = x[k]$. In Figure 6, the constant of integration `C1` is updated whenever any state is entered for the first time (line 16) this corresponds to the update of the constants of integration on the transitions in Figure 5(c).

4.4 Saturation

Due to the discrete/synchronous valuation of variables, when control remains in one location the variable may reach a value which will never satisfy the guard condition of the egress transition. This phenomena depends on the location invariant and the guard and may not happen frequently. Figure 7 illustrates three separate cases.

Case 1: Consider the step size δ as 1, the invariant at location t_1 in Figure 7(a) is $x \leq 120$ and the jump condition is $A \wedge x > 100$. According to the non-deterministic semantics of hybrid automata, the value of x, as it leaves location t_1, can be in the interval $(100, 120]$ when the signal A is present. The value of x is plotted against time in Figure 7(d) (top). In the synchronous approximation (shown as points), we have two scenarios. (1) If the transition is taken at the 5^{th} tick (time equals 4 seconds), the value of x as it leaves location t_1 is 111.2, respectively. (2) Otherwise, at the 6^{th} tick (time equals 5) the value of x is saturated to 120 and is forced to exit location t_1. Note that the trace due to the synchronous approximation is different but, still a valid trace of the HA.

This example satisfies the restrictions proposed by Alur et al. [23] and will be accepted for code generation. Here, the duration between the occurrence of the jump condition evaluating to true and the occurrence of the invariant evaluating to false is longer than the tick length (sampling period). This restriction ensures that there is always at least one valid state where a discrete transition can be taken. However, the above restriction may not always hold as illustrated in the following case.

Case 2: Given the step size as 1, Figure 7(b) illustrates a case where we must saturate. Figure 7(d) shows the value of y increasing steadily to 50 and then remaining at 50. In contrast to Case 1, for location t_3 the duration between the guard ($y = 50$) evaluating to true and the occurrence of the invariant evaluating to false ($y < 50$) is zero. This violates the restrictions proposed in [23] and this HA would not be accepted in their framework. In Simulink®, this HA may not be simulated correctly because of the equality check. In fact, for two out of seven benchmarking examples, we observed incorrect behaviour in Simulink® (see Table 2 in Section 5).

In the synchronous approximation, we accept this HA because of the proposed *saturation* technique. In our approach, during the 4^{th} tick the value of y is below 50 and the guard condition is not satisfied. During the 5^{th} tick the value of y is above 50 and without saturation this leads to an unreachable state because neither the invariant ($y < 50$) of t_3 nor the guard ($y = 50$) are valid. To address this problem, at the start of the 5^{th} tick, we saturate the value of y to exactly 50. Hence, a value greater than 50 is not observed between ticks 4 and 5. By using saturation, unlike [23], we are agnostic to the step size (sampling period).

Case 3: Similar to Case 1, Figure 7(c) illustrates a decreasing function. Once again, the duration between the occurrence of the guard evaluating to true and the occurrence of the invariant evaluating to false is longer than the step size (sampling period). This restriction ensures that there always exists at least one valid state where a discrete transition can be taken. This HA is accepted for code generation by [23] and by our tool Piha.

Definition 10 formalises this *saturation* technique.

Definition 10. *In any discrete time instant k, when execution makes a discrete switch from state l to l' in the SWA, the valuation of all continuous variables $X[k]$ either satisfy the guard condition $g \in CV(X[k])$ or are set to a suitable value such that g is satisfied. This is termed as* saturation.

According to the above definition, during the execution of the backend code, we must decide dynamically when to saturate and also must decide on the correct value. This decision is based on the following Lemma that ensures that the value of $x[k]$ that satisfies the guard always exists in the current discrete step.

Lemma 1. *It is always possible to uniquely determine the saturation value for any continuous variable at time instant k when the state (location) switch from l to l' is to be taken in a SWA.*

Proof. The proof of this lemma follows from the following observations.

- Observation 1: All witness functions $x(t)$ for any $x \in X$ are monotonic in every location (WHA requirement).
- Observation 2: All witness functions $x(t)$ for any $x \in X$ are continuous as they are differentiable in any interval.
- Observation 3: Given the above two observations, the saturation value for any variable x always exists in the time interval $[(k-1) \times \delta, k \times \delta]$ when the location switch happens at instant $k \times \delta$.

□

4.5 Composition of SHAs

This section describes the modular compilation of SHAs using the composition rules defined in Definition 7. Let us consider a part of the water tank and gas burner HAs as shown in Figure 8(a). The equivalent partial SHAs for these two HAs are shown in Figure 8(b). As described previously, in Section 4.2, the ODEs (flow predicates) in each location have been replaced with their individual witness functions. We can furthermore compile these SHAs into individual SWAs as shown in Figure 8(c). The composition rules defined in Definition 7 are applied on these resultant SWAs to generate a single SWA shown in Figure 8(d).

The pseudo-code used to compose two SWAs is shown in Algorithm 3. The algorithm takes as input two SWAs: Q_1 and Q_2, respectively. In case of N SWAs, the composition procedure is applied recursively. As the very first step, the algorithm builds the cross product of all states in the constituent SWAs Q_1 and Q_2 (line 1). Next, the

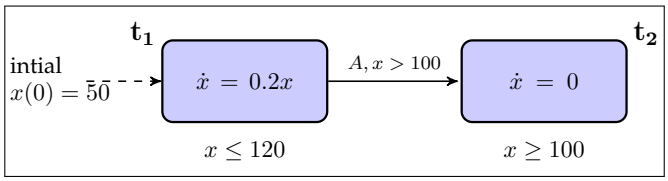

(a) Case 1: an increasing function and it does not need saturation

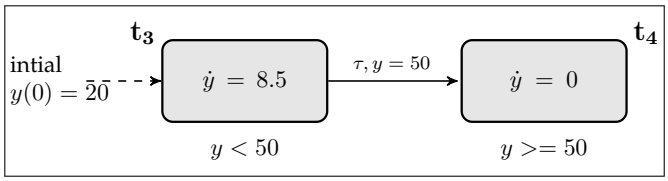

(b) Case 2: due to equality there is a need for saturation

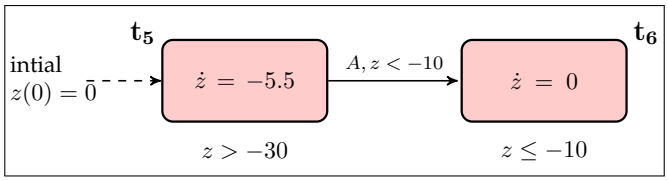

(c) Case 3: a decreasing function and it does not need saturation

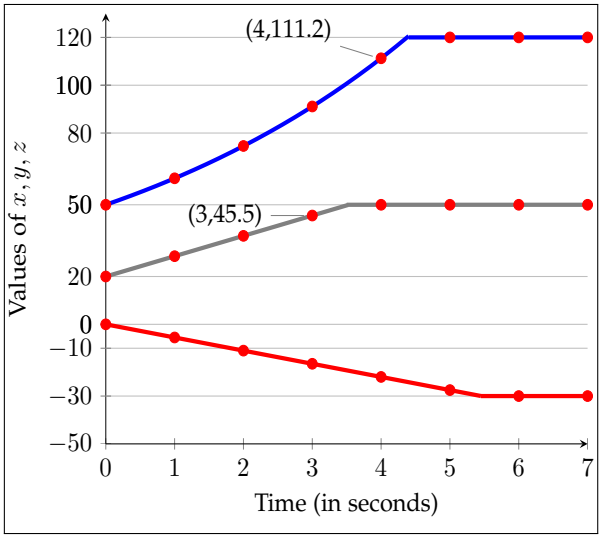

(d) The behaviours of the three example HAs are depicted using solid lines. Our synchronous approximations are depicted using dashed lines. Each tick is one second long.

Fig. 7. The need for saturation depends on the location invariant, the guard in HA, and the step size. Out of the three cases, only Case 2 requires saturation, see Figure 7(d).

Algorithm 3 The pseudo-code used to compose two SWAs

Input: SWA Q_1, Q_2
1: $Q_c \leftarrow Q_1 \times Q_2$
// Apply `Rule Intra-Intra`
2: **for all** $q \in Q_c$ **do**
3: build_self_transition(q)
4: **end for**
// Apply `Rule Intra-Inter`
5: **for all** $q_1 \in Q_c$ **do**
6: $(l_1, l_2) \leftarrow q_1$
7: **for all** $q_2 \in Q_c \wedge q_1 \neq q_2$ **do**
8: $(l_1', l_2') \leftarrow q_2$
9: **if** $l_1 = l_1' \wedge l_2 \neq l_2'$ **then**
10: build_transition(q_1, q_2)
11: **end if**
12: **end for**
13: **end for**
// `Rule Inter-Intra` is similar to the `Rule Intra-Inter`

intra-location and inter-location rules from Definition 7 are applied sequentially.

The very first rule that we apply is the intra-location transition rule (`Rule Intra-Intra`). The application of this rule simply requires one to iterate through each state in the product set Q_c, building self-transitions on states. The example of the application of `Rule Intra-Intra` is shown in Figure 8(d). Given state (t_2, b_4) in the resultant SWA; a self-transition from (t_2, b_4) to (t_2, b_4) is introduced. This transition indicates the simultaneous evolution of the ODEs from individual states t_2 and b_4 from the constituent SWAs Q_1 and Q_2, respectively. As seen from Figure 8(d), the continuous variables x and y evolve together according to their individual witness functions from the self-transitions on states t_2 and b_4 in Figure 8(c), provided that the *conjunction* of the individual guards holds.

Next, we apply the three inter-location rules sequentially. Algorithm 3 only shows the application of `Rule Intra-Inter`, since other rules can be trivially derived from this rule. The algorithm traverses through each state of the the product set Q_c. Upon visiting a state q_1, the state label is first decomposed into its constituent parts (line 6). For the running example, given $q_1 = (t_2, b_4)$, line 6, gives $l_1 = t_2$ and $l_2 = b_4$, respectively. Next, we iterate through all the states in Q_c other than state q_1, again decomposing the state label into its constituent location names, l_1' and l_2', respectively at line 8. `Rule Intra-Inter` states that the second SWA Q_2 makes an inter-location transition. SWA Q_1 on the other hand is forced to make a transition, such that the destination state after the transition has the same location label as the source transition state. Hence, the algorithm builds transitions from state q_1 to any state q_2 such that the constituent label l_1 of state q_1 and l_1' of state q_2 are the same, but l_2 and l_2' are different. The result of application of such a rule is shown in Figure 8(c), as transition $(t_2b_4, \frac{\overline{ON} \wedge \overline{OFF} \wedge (0 \leq x[k] \leq 100) \wedge (y[k] = \frac{1}{10})}{y[k+1]=y[k], x[0]=x[k], OFF, k=0}, t_2b_1)$. The second SWA Q_2 takes a discrete transition, forcing the first one to also take a discrete transition. This transition is only to a state where the location of the first SWA (Q_1) does not change. The guards on this transition are a conjunction of the set of individual guards. We add a special update $x[0] = x[k]$, which carries the value of the continuous variable x. Note that the update $x[0] = x[k]$ is the consequence of $i_1' = v_1$ in `Rule Intra-Inter`.

5 Results

We compare the efficacy of the proposed approach (our tool Piha) with Simulink. In particular we compare the

performance of these two tools relative to execution time and code size. For the purposes of this comparison, we use the seven benchmarks presented in Table 1. These benchmarks span across different application domains such as medical, physics, and industrial automation, illustrating the diversity of the proposed approach. In our setting we mean IOHA [29] when we state HA. IOHA enable the modelling of the plant and the controller separately.

As depicted in column one of Table 1, four out of the seven benchmarks are described using a single HA and the remaining three examples are described using more than one HA. The table also presents the number of locations (#L) in each hybrid automata. For example, (2,3) denotes that the Train Gate Control (TG) benchmark is described using a HA with two locations and a second HA with three locations. More details about the benchmarks and their implementation in Piha and Simulink are available online [41].

5.1 Experimental set-up

The following steps are considered in order to achieve a fair comparison between Piha and Simulink.

Solver To reflect the synchronous execution model, we used a discrete solver with a fixed step in Simulink. The Discrete-Time Integrator block is configured to use the Forward Euler method. Other methods such as Backward Euler and Trapezoidal resulted in an "algebraic loop" error and we did not pursue a solution to this.

Step size For all benchmarks the step size in Simulink is fixed to 0.01 seconds. Also the same step size is used in Piha, $\delta = 0.01$ seconds.

Time All benchmarks were simulated in Simulink for $100,000$ seconds of simulation time. Based on a step size of 0.01 seconds, in Piha this translates to 10 million ticks.

The experiments are executed on an Intel i7 processor with 16 GB RAM running the Windows 7 operating system.

5.2 Evaluation

For all the benchmarks, the executable for the Simulink models are generated using the in-built C code generator. It automatically generates equivalent C code and compiles it to produce an executable. Similarly, Piha generates equivalent C code and generates an executable using GCC. The execution time and the code size of the generated executables are reported below.

Execution time: Figure 9(a) shows that for all benchmarks, the execution times of Piha are significantly shorter than Simulink. On average, Piha is 3.9 times faster than Simulink. The most significant difference between the tools is observed for the most complex example (most locations), the water tank heating system (WH). For this example, the execution time of Piha is 7.3 times faster than Simulink.

Code size: Figure 9(a) shows that for all benchmarks, the code size of Piha is significantly smaller than Simulink. On average, the generated code is 40% smaller than the Simulink code. The most significant difference is 47% which is observed for the Thermostat (TS) example. In general, Simulink is a more feature rich tool and there may be some overheads during code generation where numerical solvers are linked during compilation.

TABLE 2
What benchmarks can be accepted based on the restrictions of the tool in [23] and our tool Piha.

Benchmarks	Tool in [23]	Piha
Thermostat (TS)	Yes	Yes
Switch Tank (ST)	Yes	Yes
Heart Cell (HC)	Yes	Yes
Train Brake control (TB)	Yes	Yes
Water Heating system (WH)	No	Yes
Train Gate control(TG)	No	Yes
Nuclear Plant control (NP)	Yes	Yes

In summary, on average Piha is faster (in execution time) than Simulink by a factor of 3.9 times and the generated code is 40% smaller than the Simulink code. In the future, we will extend our benchmark suite with more complex examples and quantitatively compare with Ptolemy and Zélus.

Finally, in Table 2 we qualitatively compare the acceptability of the benchmarks between our tool and the tool based on [23]. As discussed earlier in the introduction and Section 4.4, the tool in [23] requires that "a given mode and the corresponding guard of any switch must overlap for a duration that is greater than the sampling period" [23]. Due to this restriction, it is not possible to accept any HA that has a guard condition that checks for equality. For the running example (Figure 2(c)), the guard condition $x = 100$ checks if the water has reached the boiling point. Similarly, in the train gate control benchmark, there is a guard condition that checks if the gate height is exactly equal to 10 meters. In summary, due to our saturation function, we can accept a larger set of benchmarks for code generation than the tool in [23].

6 Conclusions

Hybrid automata (HA) [4] is a very well known framework for the modelling and verification of Cyber-physical Systems (CPS) [2]. The primary focus of the current work is the emulation of controllers using plant models that provide real-time closed-loop response, while validating the controllers in a CPS. We term such plant models as plant-on-a-chip (PoC).

HA was initially developed for the simulation / verification of CPS where both the plant and the controller are considered as a single HA. We seek to validate controllers using emulation of the plant. Variants such as input / output HA [29] (IOHA) are proposed to specify both the plant and the adjoining controller and this work has recently been used for the validation of controllers such as pacemakers [28]. The majority of controller validation approaches use tools such as Simulink for the modelling and code generation of the plant to validate the controller. There are well known semantic limitations of using such tools for validation. Tools such as Ptolemy [13] and Zélus [12], on the other hand, are founded on formal semantics. However, they have limitations for emulating plants due to the dynamic interaction with numerical solvers. There have also been some prior work on automatic code generation from

TABLE 1
Benchmark descriptions

	Benchmarks	Domain	# L	Description
Single HA	Thermostat (TS)	Physics [42]	2	Heats a room to keep it warm.
	Switch Tank (ST)	Physics [43]	2	A hose adding water by switching between two leaky tanks.
	Heart Cell (HC)	Biology [44]	4	Captures the electrical behaviour of a cardiac cell.
	Train Brake control (TB)	Industrial automation [45]	2	Maintains a train's speed between the upper and lower limit.
Two HAs	Water Heating system (WH)	Physics [3]	4 , 4	Models the heating of water (see Figure 2)
	Train Gate control(TG)	Industrial automation [46]	2, 3	Models the behaviour of a gate at a rail road crossing.
	Nuclear Plant control (NP)	Industrial automation [2]	3, 3	Switches between two fuel rods to avoid a meltdown.

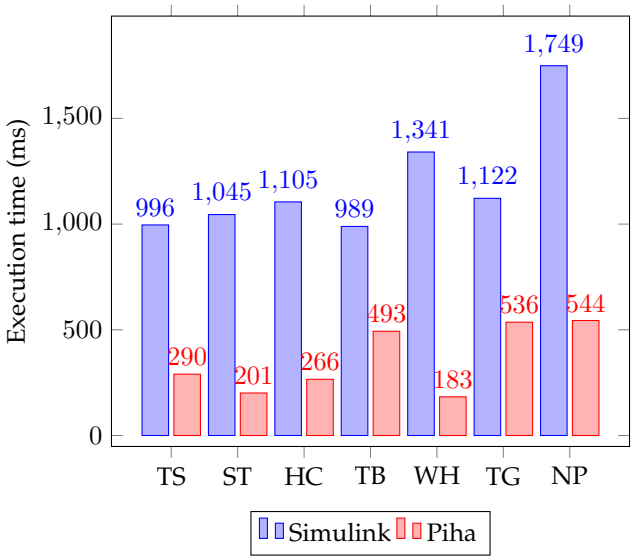

(a) Execution time (ms)

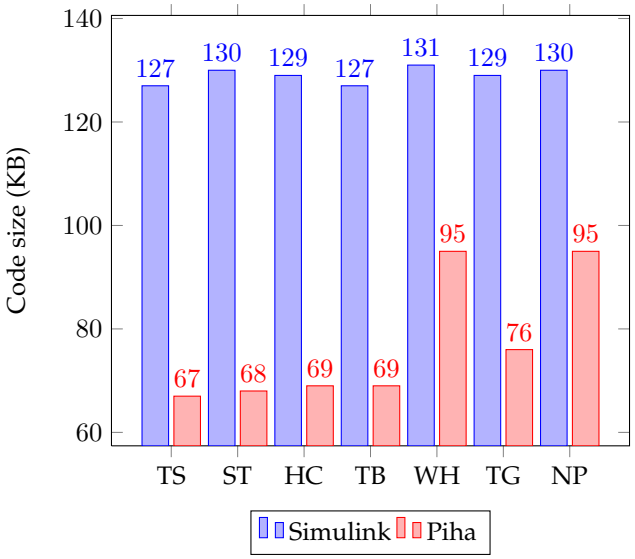

(b) Code size (KB)

Fig. 9. Comparing the execution time (in ms) and code size (in KB) of Simulnk and Piha for the benchmarks in Table 1.

HA models [23], [24]. However, these approaches handle very restrictive set of examples.

This paper formulates the problem of emulation of CPS using automated algorithmic techniques for code generation from IOHA models. We firstly defined a set of well-formedness criteria that are specifically developed for facilitating code generation. We also propose a discrete time semantics of well-formed HA (WHA) based on the synchronous approach [16]. Based on this semantics and an approach for synchronous composition of HA models, also proposed here, we are able to perform modular compilation from a network of HA models to C code for PoC design.

Experimental validation of the proposed approach compared with Simulink reveals that the generated code is efficient both in terms of execution time and code size while avoiding the semantic ambiguities associated with Simulink. The proposed approach, thus, paves the way for the emulation of a wide range of CPS applications in automotive, medical devices, and robotics.

In the near future, we will compare Piha with Zélus and Ptolemy quantitatively. We will also pursue the decidability question of WHA.

ACKNOWLEDGEMENTS

This research was funded by the University of Auckland FRDF Postdoctoral grant No. 3707500. The authors would like to thank Théo Steger, an intern from the ENSTA Paris Tech University, for his help in developing benchmarks and testing our tool. The authors also acknowledge discussions with colleagues in Control Engineering, in particular Nitish Patel and Akshya Swain for discussions on sampled systems and synchronous compositions between plant and controllers.

REFERENCES

[1] E. A. Lee, "Cyber Physical Systems: Design Challenges," in *Proceedings of the 2008 11th IEEE Symposium on Object Oriented Real-Time Distributed Computing*, ISORC '08, (Washington, DC, USA), pp. 363–369, IEEE Computer Society, 2008.
[2] R. Alur, *Principles of Cyber-Physical Systems*. MIT Press, 2015.
[3] J.-F. Raskin, *Handbook of Networked and Embedded Control Systems*, ch. An introduction to hybrid automata, pp. 491–517. Springer, 2005.
[4] R. Alur, C. Courcoubetis, T. A. Henzinger, and P.-H. Ho, "Hybrid Automata: An Algorithmic Approach to the Specification and Verification of Hybrid Systems," in *Hybrid Systems*, (London, UK), pp. 209–229, Springer-Verlag, 1993.
[5] R. Wilhelm, J. Engblom, A. Ermedahl, N. Holsti, S. Thesing, D. Whalley, G. Bernat, C. Ferdinand, R. Heckmann, T. Mitra, F. Mueller, I. Puaut, P. Puschner, J. Staschulat, and P. Stenström, "The worst-case execution-time problem—overview of methods and survey of tools," *Trans. on Embedded Computing Systems, ACM*, vol. 7, no. 3, pp. 1–53, 2008.

[6] International Electrotechnical Commission, "IEC 61508 Functional safety of electrical / electronic / programmable electronic safety-related systems." http://www.iec.ch/functionalsafety/. last accessed 01/09/2015.
[7] International Organization for Standardization, "ISO 26262-1: Road vehicles – Functional safety." http://www.iso.org/iso/catalogue˙detail?csnumber=43464. last accessed 01/09/2015.
[8] W. K. Youn, S. B. Hong, K. R. Oh, and O. S. Ahn, "Software certification of safety-critical avionic systems: DO-178C and its impacts," *Aerospace and Electronic Systems Magazine, IEEE*, vol. 30, no. 4, pp. 4–13, 2015.
[9] "US Food and Drug Administration (FDA)." http://www.fda.gov. last accessed 29/05/2015.
[10] H. Alemzadeh, R. K. Iyer, Z. Kalbarczyk, and J. Raman, "Analysis of safety-critical computer failures in medical devices," *Security & Privacy, IEEE*, vol. 11, no. 4, pp. 14–26, 2013.
[11] R. Alur, A. Kanade, S. Ramesh, and K. Shashidhar, "Symbolic analysis for improving simulation coverage of Simulink/Stateflow models," in *Proceedings of the 8th ACM international conference on Embedded software*, pp. 89–98, ACM, 2008.
[12] T. Bourke and M. Pouzet, "Zélus: a synchronous language with ODEs," in *Proceedings of the 16th international conference on Hybrid systems: computation and control*, pp. 113–118, ACM, 2013.
[13] C. Ptolemaeus, *System Design, Modeling, and Simulation: Using Ptolemy II*. Ptolemy. org, 2014.
[14] T. Bourke, J.-L. Colaco, B. Pagano, C. Pasteur, and M. Pouzet, "A synchronous-based code generator for explicit hybrid systems languages," in *Compiler Construction* (B. Franke, ed.), vol. 9031 of *Lecture Notes in Computer Science*, pp. 69–88, Springer Berlin Heidelberg, 2015.
[15] K. N. Patel and R. H. Javeri, "A survey on emulation testbeds for mobile ad-hoc networks," *Procedia Computer Science*, vol. 45, pp. 581–591, 2015.
[16] A. Benveniste, P. Caspi, S. Edwards, N. Halbwachs, P. Le Guernic, and R. de Simone, "The synchronous languages 12 years later," *Proceedings of the IEEE*, vol. 91, pp. 64–83, Jan. 2003.
[17] C. Baier and J.-P. Katoen, *Principles of Model Checking*. The MIT Press, 2008.
[18] G. Frehse, C. Le Guernic, A. Donzé, S. Cotton, R. Ray, O. Lebeltel, R. Ripado, A. Girard, T. Dang, and O. Maler, "SpaceEx: Scalable Verification of Hybrid Systems," in *Proceedings of the 23rd International Conference on Computer Aided Verification*, CAV'11, (Berlin, Heidelberg), pp. 379–395, Springer-Verlag, 2011.
[19] S. Tripakis, C. Sofronis, P. Caspi, and A. Curic, "Translating discrete-time Simulink to Lustre," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 4, no. 4, pp. 779–818, 2005.
[20] S. Andalam, P. S. Roop, A. Girault, and C. Traulsen, "A predictable framework for safety-critical embedded systems," *IEEE Transactions on Computers*, vol. 63, no. 7, pp. 1600–1612, 2014.
[21] "SCADE Tools." http://www.esterel-technologies.com/. last accessed - 19.6.15.
[22] X. Fornari, "Understanding how SCADE suite KCG generates safe C code," *White paper, Esterel Technologies*, 2010.
[23] R. Alur, F. Ivancic, J. Kim, I. Lee, and O. Sokolsky, "Generating embedded software from hierarchical hybrid models," *ACM SIGPLAN Notices*, vol. 38, no. 7, pp. 171–182, 2003.
[24] J. Kim and I. Lee, "Modular code generation from hybrid automata based on data dependency," in *Real-Time and Embedded Technology and Applications Symposium, 2003. Proceedings. The 9th IEEE*, pp. 160–168, IEEE, 2003.
[25] D. Bresolin, L. Di Guglielmo, L. Geretti, and T. Villa, "Correct-by-construction code generation from hybrid automata specification," in *International Wireless Communications and Mobile Computing Conference (IWCMC)*, pp. 1660–1665, IEEE, July 2011.
[26] R. Alur and D. L. Dill, "A theory of timed automata," *Theoretical computer science*, vol. 126, no. 2, pp. 183–235, 1994.
[27] L. H. Yoong, P. S. Roop, Z. E. Bhatti, and M. M. Kuo, *Model-Driven Design Using IEC 61499*. Springer, 2015.
[28] T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre, "Quantitative verification of implantable cardiac pacemakers over hybrid heart models," *Information and Computation*, vol. 236, pp. 87–101, 2014.
[29] N. Lynch, R. Segala, and F. Vaandrager, "Hybrid I/O automata," *Information and computation*, vol. 185, no. 1, pp. 105–157, 2003.
[30] C. Brooks, E. A. Lee, D. Lorenzetti, T. S. Nouidui, and M. Wetter, "CyPhySim: A Cyber-physical Systems Simulator," in *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, HSCC '15, (New York, NY, USA), pp. 301–302, ACM, 2015.
[31] D. Harel and A. Pnueli, "On the Development of Reactive Systems," in *Logics and Models of Concurrent Systems* (K. Apt, ed.), NATO ASI Series, Vol. F-13, (La Colle-sur-Loup, France), pp. 477–498, Springer-Verlag, 1985.
[32] K. Ogata, *Modern control engineering*. Boston : Prentice-Hall, 2010.
[33] H. Carlsson, B. Svensson, F. Danielsson, and B. Lennartson, "Methods for reliable simulation-based PLC code verification," *Industrial Informatics, IEEE Transactions on*, vol. 8, no. 2, pp. 267–278, 2012.
[34] L. Ju, B. K. Huynh, A. Roychoudhury, and S. Chakraborty, "Performance Debugging of Esterel Specifications," in *IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS)*, (Atlanta), pp. 173–178, ACM, October 2008.
[35] P. S. Roop, S. Andalam, R. von Hanxleden, S. Yuan, and C. Traulsen, "Tight WCRT Analysis of Synchronous C Programs," in *Proceedings of the international conference on Compilers, architecture, and synthesis for embedded systems (CASES)*, (Grenoble), pp. 205–214, ACM, October 2009.
[36] J. J. Wang, P. S. Roop, and S. Andalam, "ILPc: A Novel Approach for Scalable Timing Analysis of Synchronous Programs," in *International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (CASES)*, Oct. 2013.
[37] G. Berry, "The foundations of Esterel.," in *Proof, language, and interaction*, pp. 425–454, 2000.
[38] N. Halbwachs, P. Caspi, P. Raymond, and D. Pilaud, "The synchronous data flow programming language LUSTRE," *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1305–1320, 1991.
[39] P. LeGuernic, T. Gautier, M. Le Borgne, and C. Le Maire, "Programming real-time applications with SIGNAL," *Proceedings of the IEEE*, vol. 79, no. 9, pp. 1321–1336, 1991.
[40] W. Rudin, *Real and complex analysis*. Tata McGraw-Hill Education, 1987.
[41] "Piha Benchmarks." https://github.com/PRETgroup/PihaBenchmarks. last accessed - 15.08.2015.
[42] H. Joao Pedro, "How to describe a hybrid system? Formal models for hybrid system." University of California at Santa Barbara, Course ECE229, Lecture 2. http://www.ece.ucsb.edu/˜hespanha/ece229/Lectures/Lecture2.pdf, 2005.
[43] J. Lygeros, C. Tomlin, and S. Sastry, "Hybrid systems: Modeling, analysis and control." University of California, Berkeley, Course EE291e. http://inst.cs.berkeley.edu/˜ee291e/sp09/handouts/book.pdf, 2009.
[44] T. Chen, M. Diciolla, M. Kwiatkowska, and A. Mereacre", "Quantitative verification of implantable cardiac pacemakers over hybrid heart models ," *Information and Computation*, vol. 236, pp. 87 – 101, 2014.
[45] A. Platzer, "Logical analysis of hybrid systems," in *Descriptional Complexity of Formal Systems*, pp. 43–49, Springer, 2012.
[46] C. Brennon and E. Joshua, "Hybrid Systems." Tufts University, Course EE194, Lecture. http://www.eecs.tufts.edu/˜khan/Courses/Spring2013/EE194/Lecs/ Hybrid_ Systems_Presentation_Elliott_Costello.pdf, 2013.

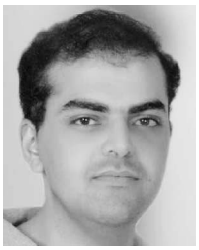

Avinash Malik Avinash Malik is a Lecturer at the University of Auckland, New Zealand. His main research interest lies in programming languages for multicore and distributed systems and their formal semantics/ compilation. He has worked at organizations such as INRIA in France, Trinity College Dublin, IBM research Ireland, and IBM Watson on design and compilation of programming languages. He holds B.E. and PhD degrees from the University of Auckland.

Partha S Roop Partha received PhD degree in computer science (software engineering) from the University of New south Wales, Sydney, Australia, in 2001. He is currently an Associate Professor and is the Director of the Computer Systems Engineering Program with the Department of Electrical and Computer Engineering, University of Auckland, New Zealand. Partha is an associated team member of the SPADES team INRIA, Rhone-Alpes and held a visiting positions in CAU Kiel and Iowa state University. His research interests include the design and verification of embedded systems. In particular, he is developing techniques for the design of embedded applications in automotive, robotics, and intelligent transportation systems that meet functional safety standards.

Sidharta Andalam Sidharta Andalam received his Ph.D from University of Auckland in 2013, where his thesis focused on developing a predictable platform for safety-critical systems. He is currently a research fellow in embedded systems at the University of Auckland, New Zealand. His principle research interests are on design, implementation and analysis of safety-critical applications. He has worked at TUM CREATE, Singapore exploring safety-critical applications in automotive domain.

Eugene Yip Eugene Yip received the B.E. (Hons) and Ph.D degrees in electrical and computer systems engineering from the University of Auckland, New Zealand. At the start of 2015, he joined the Heart-on-FPGA research group in the Department of Electrical and Computer Engineering, University of Auckland, as a research assistant working on the real-time modeling of cardiac electrical activity. In the middle of 2015, he joined the SWT research group at the University of Bamberg as a research assistant working on synchronous languages. His current research interests include synchronous mixed-criticality systems, parallel programming, static timing analysis, formal methods for organ modeling, and biomedical devices.

Mark Trew Mark Trew attained a Bachelor of Engineering in Engineering Science in 1992 and a PhD Engineering Science in 1999, both from the University of Auckland. He is currently a Senior Research Fellow at the Auckland Bioengineering Institute. Mark constructs computer models and analysis tools for interpreting and understanding detailed images of cardiac tissue and cardiac electrical activity.

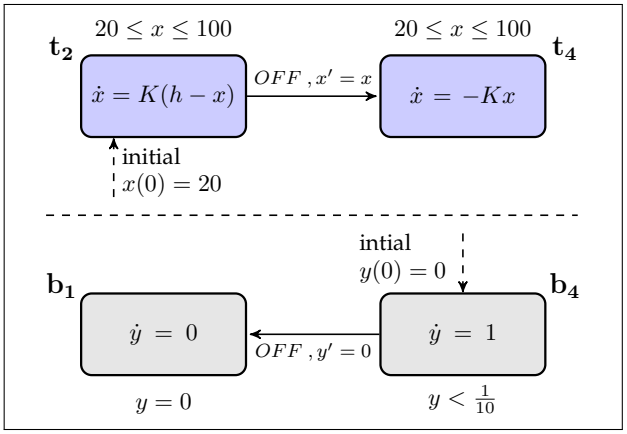

(a) Partial hybrid automata of the water tank (top) and gas burner (bottom), reproduced from Figure 2.

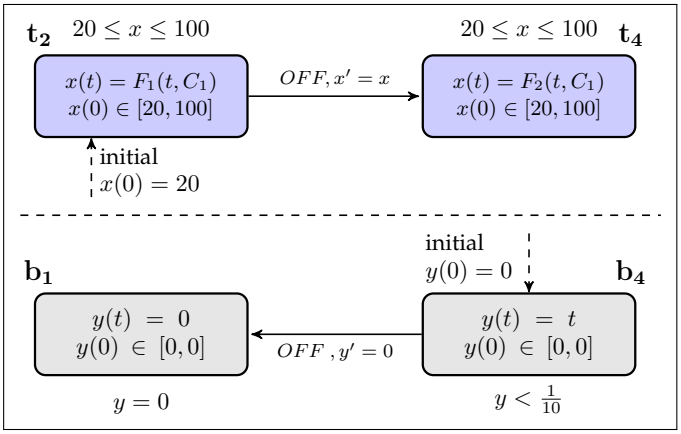

(b) SHA of water tank (top) and gas burner(bottom)

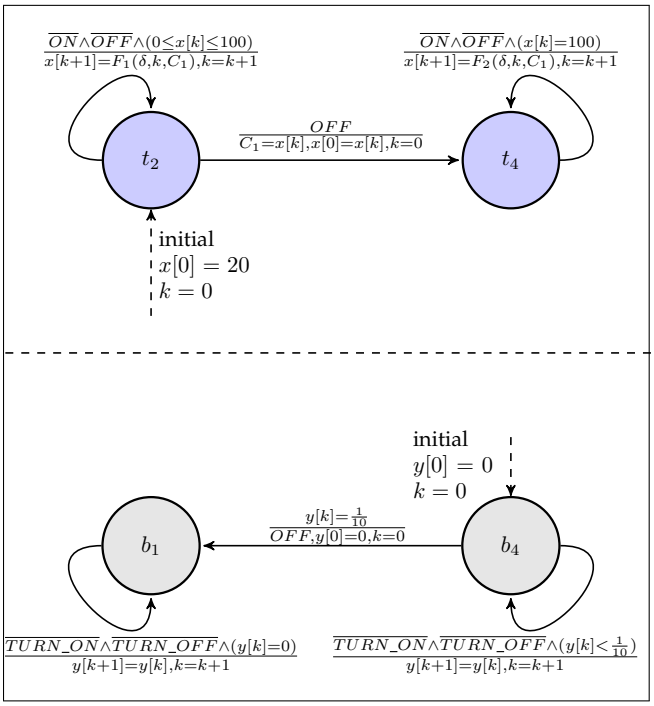
$\frac{\overline{ON} \wedge \overline{OFF} \wedge (0 \leq x[k] \leq 100)}{x[k+1] = F_1(\delta, k, C_1), k = k+1}$
$\frac{\overline{ON} \wedge \overline{OFF} \wedge (x[k] = 100)}{x[k+1] = F_2(\delta, k, C_1), k = k+1}$
t_2
$\frac{OFF}{C_1 = x[k], x[0] = x[k], k = 0}$
t_4
initial
$x[0] = 20$
$k = 0$
initial
$y[0] = 0$
$k = 0$
b_1
$\frac{y[k] = \frac{1}{10}}{OFF, y[0] = 0, k = 0}$
b_4
$\frac{\overline{TURN\_ON} \wedge \overline{TURN\_OFF} \wedge (y[k] = 0)}{y[k+1] = y[k], k = k+1}$
$\frac{\overline{TURN\_ON} \wedge \overline{TURN\_OFF} \wedge (y[k] < \frac{1}{10})}{y[k+1] = y[k], k = k+1}$

(c) Synchronous Witness Automata (SWA) of water tank and gas burner.

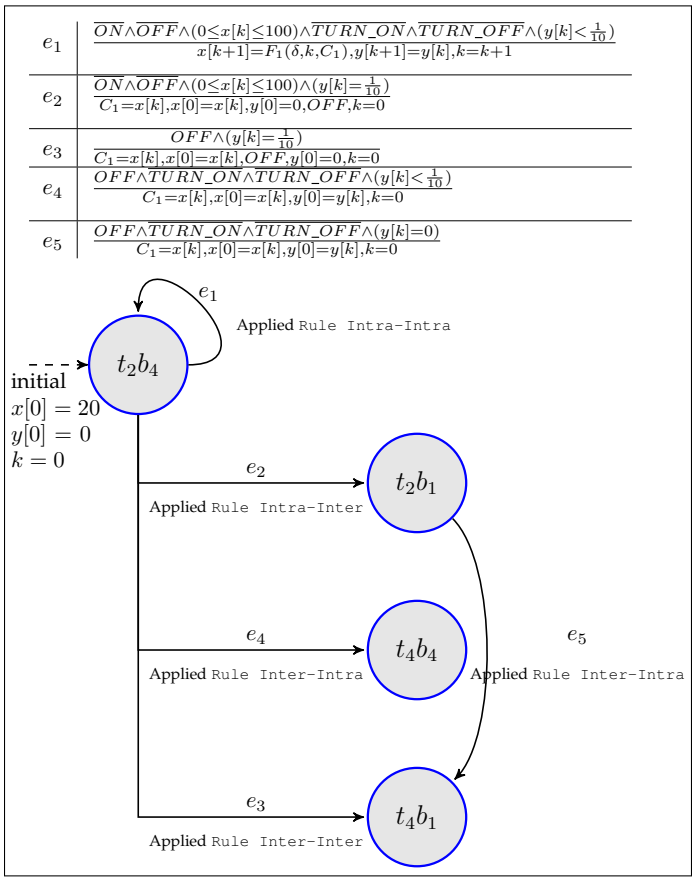
e_1 $\frac{\overline{ON} \wedge \overline{OFF} \wedge (0 \leq x[k] \leq 100) \wedge \overline{TURN\_ON} \wedge \overline{TURN\_OFF} \wedge (y[k] < \frac{1}{10})}{x[k+1] = F_1(\delta, k, C_1), y[k+1] = y[k], k = k+1}$
e_2 $\frac{\overline{ON} \wedge \overline{OFF} \wedge (0 \leq x[k] \leq 100) \wedge (y[k] = \frac{1}{10})}{C_1 = x[k], x[0] = x[k], y[0] = 0, OFF, k = 0}$
e_3 $\frac{OFF \wedge (y[k] = \frac{1}{10})}{C_1 = x[k], x[0] = x[k], OFF, y[0] = 0, k = 0}$
e_4 $\frac{OFF \wedge \overline{TURN\_ON} \wedge \overline{TURN\_OFF} \wedge (y[k] < \frac{1}{10})}{C_1 = x[k], x[0] = x[k], y[0] = y[k], k = 0}$
e_5 $\frac{OFF \wedge \overline{TURN\_ON} \wedge \overline{TURN\_OFF} \wedge (y[k] = 0)}{C_1 = x[k], x[0] = x[k], y[0] = y[k], k = 0}$
e_1
Applied Rule Intra-Intra
t_2b_4
initial
$x[0] = 20$
$y[0] = 0$
$k = 0$
e_2
Applied Rule Intra-Inter
t_2b_1
e_4
Applied Rule Inter-Intra
t_4b_4
e_5
Applied Rule Inter-Intra
e_3
Applied Rule Inter-Inter
t_4b_1

(d) Parallel composition of the SWAs of water tank and gas burner.

Fig. 8. Composition of HAs in our framework